*Simplifying Cyber Security since 2016*

# HACKERCOOL

October 2025 Edition 8 Issue 10

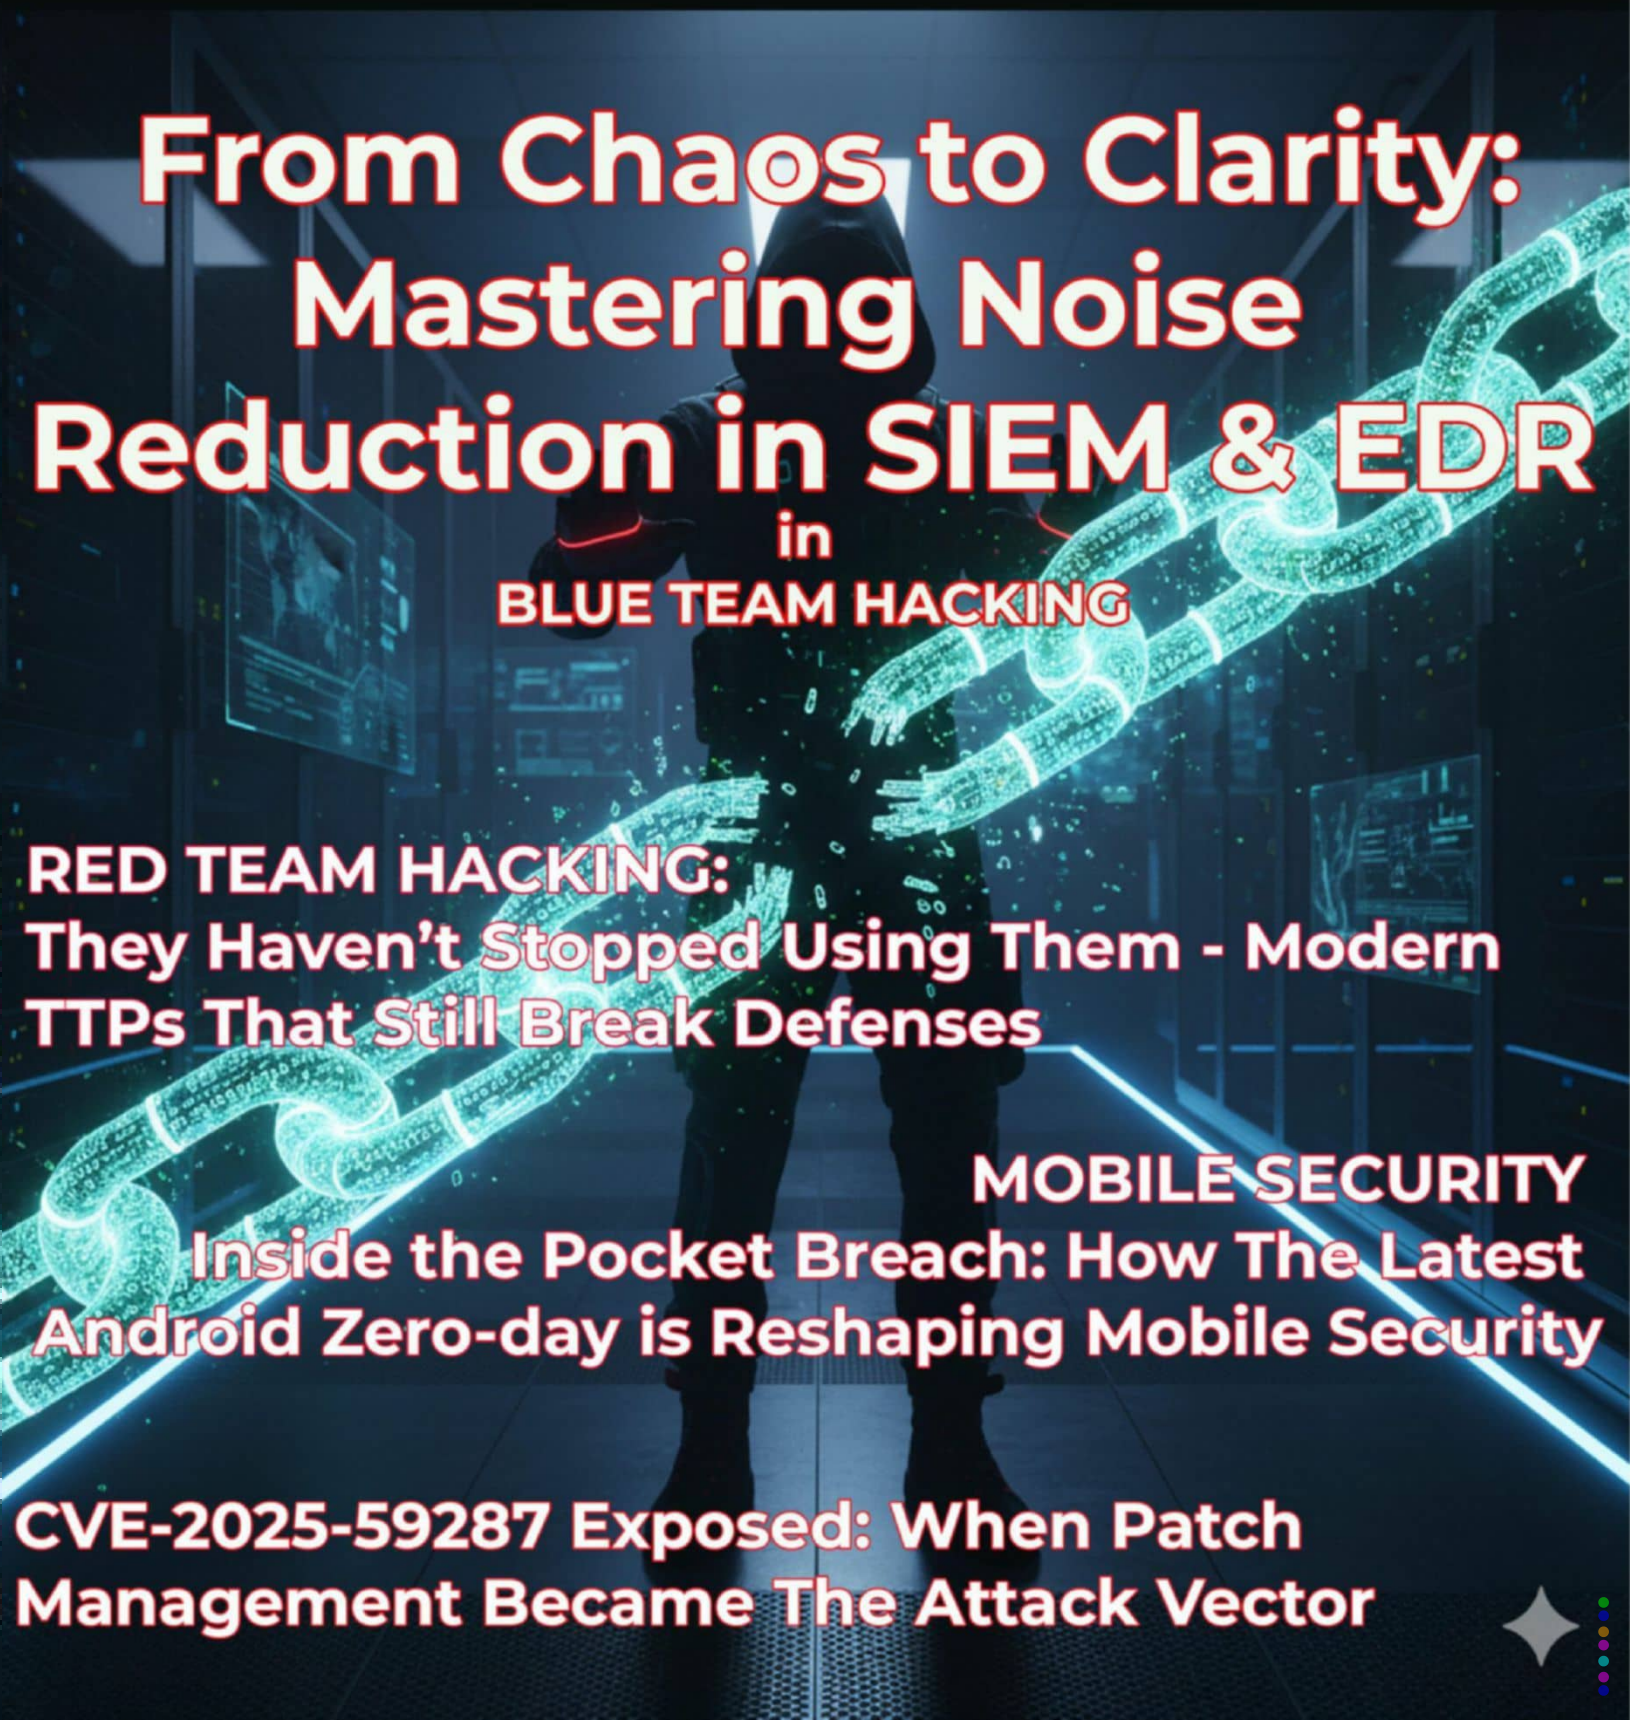

## From Chaos to Clarity: Mastering Noise Reduction in SIEM & EDR in BLUE TEAM HACKING

## RED TEAM HACKING: They Haven't Stopped Using Them - Modern TTPs That Still Break Defenses

## MOBILE SECURITY Inside the Pocket Breach: How The Latest Android Zero-day is Reshaping Mobile Security

## CVE-2025-59287 Exposed: When Patch Management Became The Attack Vector

Copyright © 2025 Hackercool CyberSecurity (OPC) Pvt Ltd

All rights reserved. No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publishe -r, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law. For permission requests, write to the publisher, addressed “Attention: Permissions Coordinator,” at the address below.

Any references to historical events, real people, or real places are used fictitiously. Na -mes, characters, and places are products of the author’s imagination.

Hackercool Cybersecurity (OPC) Pvt Ltd.
Banjara Hills, Hyderabad 500034
Telangana, India.

Website :
www.hackercoolmagazine.com

Email Address :
admin@hackercoolmagazine.com



HACKERCOOL
Simplifying Cybersecurity

## DISCLAIMER

Information provided in this Magazine is strictly for educational purpose only. Please don't misuse this knowledge to hack into devices or networks without taking relevant permissions. The Magazine will not take any responsibility for misuse of this information.

Then you will know the truth and the truth will set you free.
John 8:32

# Editor's Note

*Edition 8 Issue 10*

In cybersecurity, every month feels like its own era. Threats evolve, tools mature and the line between offense and defense blurs just a little more. In this Issue, we dive head-first into that shifting landscape, bringing you a cura -ted mix of red team tactics, blue team strategy, vulnerability breakdowns, platform updates and the growing concerns around mobile and online security.

We open with “Inside the Kill Chain: Modern TTPs That Still Work”, a red-team deep dive into the techniques that continue to evade detection even in 2025. While the industry talks endlessly about AI-driven defense, this feature reminds us that fundamentals that are executed well still remain painfully effective.

Balancing offense with defense, we bring you our blue team anchor “The Art of Noise Reduction” which explores how SOC teams can reclaim clarity by tuning SIEM and EDR systems for signal over fatigue. As alert vol -umes explode, noise reduction is no longer optimization—it's survival.

This Issue also puts a strong spotlight on latest discovered vulnerabilities, breaking down four major CVEs shaping enterprise risk this year. From Lan Scope’s endpoint flaw (CVE-2025-61932) to UniFi Access’s physical-security bypass (CVE-2025-52665), from Adobe Commerce and Magento’s critical exposure (CVE-2025-54236) to the WSUS weakness in Windows Servers (CVE-2025-59287). Each explainer demystifies the technical details so that n-ew practitioners can build real-world understanding.

For defenders relying on visibility tools, we cover Security Onion 2.4.180, walking through what’s new and what matters in live enterprise environment -s. And because the threat is now always in our pocket, our mobile section analyzes the latest Android zero-day reshaping the way we think about smar-tphone security.

Finally, we close with a look toward consumer tech and privacy. With OpenAI’s Atlas browser, convenience and seamlessness come at a price and this article examines what the marketing gloss doesn’t tell you. As always, thank you for joining us. Whether you’re learning security, defending networks, testing boundaries or simply staying informed, this Issue brings the clarity and depth that the modern threat landscape demands. Stay sharp, stay curious—and stay secure.

## Contact us:

Mail: admin@hackercoolmagazine.com
WhatsApp/Telegram: +919505658443



# INSIDE

See what our Hackercool Magazine's October 2025 Issue has in store for you.

1. Red Team Hacking:
Inside the Kill Chain: Modern TTPs That Still Work.

2. Vulnerability for beginners:
Inside CVE-2025-61932: How a LanScope Vulnerability Threatens EndPoint Security

3. What's New:
Security Onion 2.4.180

4. Vulnerability for beginners:
CVE-2025-52665: The UniFi Access Flaw Exposing Physical Security Controls

5. Blue Team Hacking:
The Art Of Noise Reduction: Tuning SIEM And EDR for Signal over Alert Fatigue

6. Vulnerability for beginners:
CVE-2025-54236: The Critical Adobe Commerce and Magento Open-source Vulne -rability Exposing Online Stores

7. Mobile Security:
Inside The Pocket Breach: How The Latest Android Zero-day is Reshaping Mobile Security

8. Vulnerability for beginners:
Inside CVE-2025-59287: How A WSUS Flaw Opens The Door to Update Hijacking

9. Online Security:
OpenAI's Atlas Browser Provides Ultimate Convenience. But The Glossy Marketin-g Masks Safety Risks.

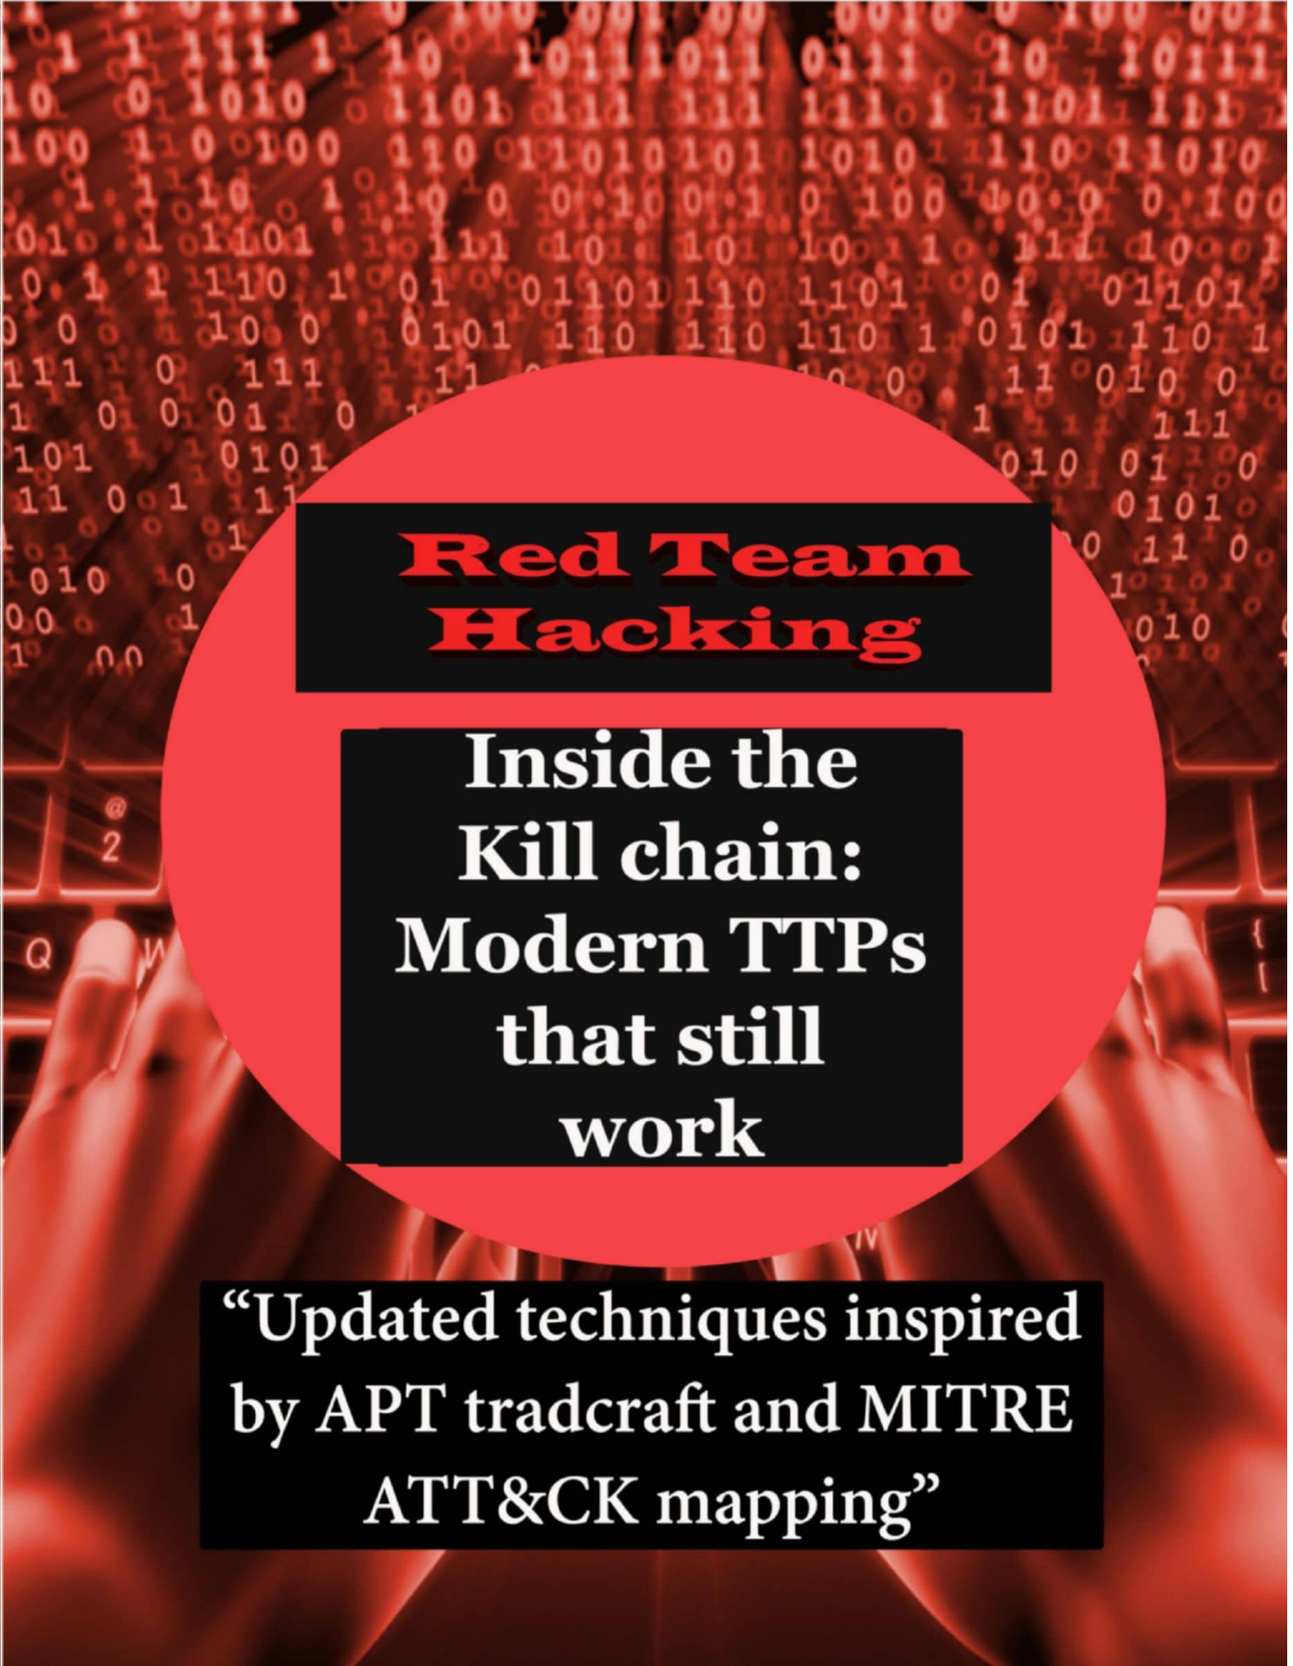


Every few years, the cybersecurity industry declares that the “kill chain” model is outdated. Yet, despite new frameworks, buzzwords and the rise of AI-driven defenses, the kill chain survives. A cyber kill chain is a framework for identifying and breaking down different stages of a cyberattack. It originates from a military concept and is developed by Lockheed Martin. Attackers still follow the same essential phases in a cyber-attack: reconnaissance, initial access, execution, persistence, privilege escalation, lateral movement, command and control and finally exfiltration or impact.

Some tools and techniques change, the target of the attack may change — but the underlying tactics remain remarkably constant. What changes are the techniques and procedures that make these tactics effective in today’s environments.

In Red Team Hacking feature of this month, we take you through a tour of the modern kill chain: the TTPs (Tactics, Techniques and Procedures) that still work, how they map to MITRE ATT&CK and what defenders can realistically do to detect and disrupt them.

## Reconnaissance: The quiet stalk

Attackers always start their hacking attack where defenders rarely look — in PUBLIC. Social media, GitHub, press releases and job postings are like gold mines for adversaries. A simple LinkedIn query can reveal the technologies used by your organization. Similarly, a press release might list your third-party vendors or managed service providers. GitHub commits can leak credentials or configuration data.

Advanced Persistent Threat (APT) actors and Red Teams alike automate this reconnaissance technique with open-source tools like theHarvester, Recon-ng and SpiderFoot combined with targeted manual research. Note that they are not just identifying assets in this phase — they’re learning who’s responsible for them, how they’re maintained and where they lay exposed.

**MITRE ATT&CK Mapping:** Reconnaissance (T1592–T1595)

**Defensive Move:** Blue Teams should treat OSINT as an attack surface. R-

egularly audit what's publicly visible about your organization. Monitor exter
-nal scanning and use deception data (fake subdomains, bogus credentials) t-
o help you detect attacker's reconnaissance early.

## Initial Access: Still a Human Problem

All the information gathered in the reconnaissance stage will prove to this s
-tage. Exploiting vulnerabilities in Internet-facing systems and Social enginee
-ring are still the most popular techniques used for initial access. The most tar
-geted internet-faced systems include vulnerable VPNs, out-of-date CMS platf
-orms, Routers, Firewalls and unpatched web applications.

If there are now vulnerabilities to expose, unique and custom designed soc
-ial engineering attacks still dominate as an effective technique to gain initial
access. If there's one social engineering technique that hasn't changed in a d
-ecade, it's phishing. Email remains the easiest way into a network because p
-eople, not technology are the weakest link.



Modern phishing campaigns, however, don't rely on malware attachments. Instead, they use Adversary-in-The-Middle (AiTM) kits that steal session tok-
ens even from MFA-protected users. Attackers deploy short-lived reverse pr-
oxies, capture the victim's cookies and immediately log in with valid credent
-ials – bypassing MFA entirely.

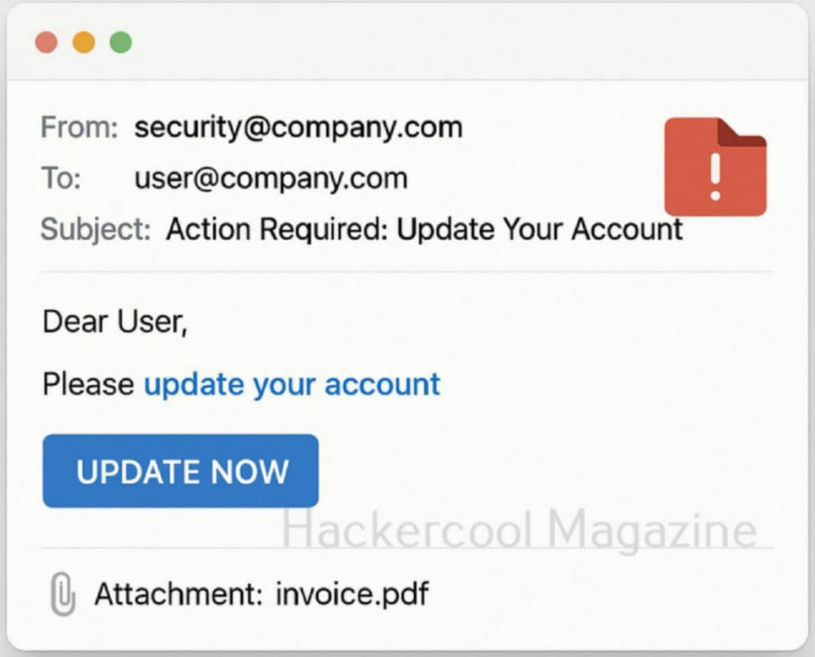


**MITRE ATT&CK Mapping:** Phishing (T1566), Exploit Public-Facing App
-lication (T1190)

**Defensive Move:** Blue Teams should deploy phishing-resistant MFA (FID
-O2, smartcards), harden external services and segment management interfa-
ces. Rapid patching and strict exposure management matter more than ever.

## Execution: Living Off the Land

Once the attacker gains initial access and is inside the target network, his g-
oal is simple. It is to execute code without raising alarms. However, the days
of dropping malware droppers are long gone.Modern Red teams and APTs

are living off the land nowadays. i.e they are using Living off The Land techniques.

PowerShell, WMI, rundll32, mshta and even certutil are being used extensively by both Red Teams and APTs to execute code on the target machines. Why? Because they're legitimate tools, signed by Microsoft and already on the system. Downloading tools from external network may raise suspicions. So, attackers' chain LOLbins together to download payloads, dump credentials or establish C2 – all while looking like normal administrative activity.

More recently, adversaries are also abusing Microsoft Teams and OneNote to deliver payloads, exploiting users' trust in internal collaboration tools.

**MITRE ATT&CK Mapping:** Command and Scripting Interpreter (T1059), Signed Binary Proxy Execution (T1218).

**Defensive Move:** Blue Teams should implement script logging (AMSI, Sysmon) and monitor for unusual command-line arguments. Application allowlisting can cripple LOTL abuse, though it requires cultural buy-in and tuning.



**Fig: Common "Living Off the Land" Binaries to Watch**

| Binary | Typical Abuse | Example Indicator |
|---|---|---|
| rundll32.exe | Executes DLL payloads | Unusual DLL paths or encoded arguments |
| mshta.exe | Runs malicious HTML apps | Network calls to external domains |
| certutil.exe | Downloads or encodes payloads | Use of -decode or remote URLs |
| wmic.exe | Remote execution | WMI process creation events |
| powershell.exe | Fileless execution | Base64-encoded command lines |

Hackercool Magazine

## Persistence: Staying Power in the Shadows

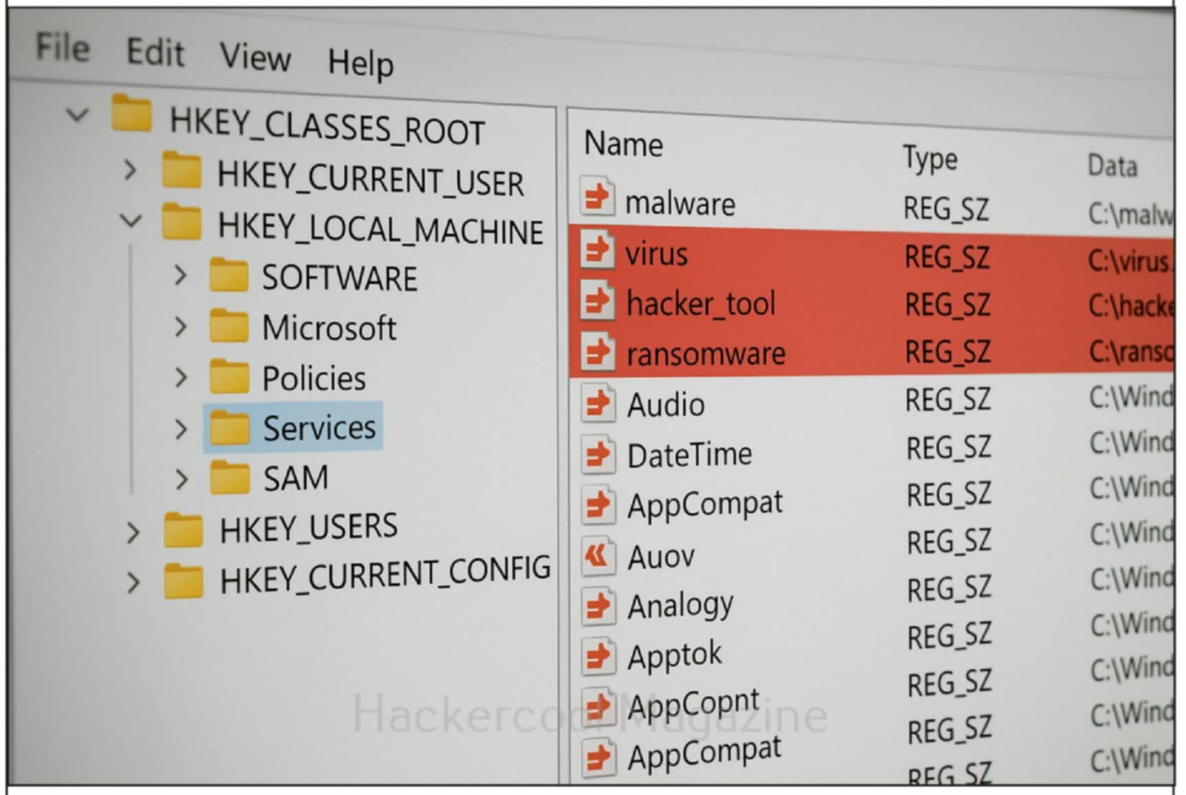

Establishing persistence is very important after gaining access. Persistence is always about subtlety, not flash. Attackers rarely need to create new services anymore, they just hijack existing ones.

Red teams now modify scheduled tasks to call malicious payloads, use regi -stry Run keys or exploit Windows Management Instrumentation (WMI) eve -nt subscriptions. In cloud environments, persistence might involve adding a service principal with elevated roles or embedding malicious logic in a CI/ CD pipeline.

**MITRE ATT&CK Mapping:** Boot or Logon Autostart (T1547), Create or Modify System Process (T1543), Cloud Account (T1136)

**Defensive Move:** Blue Teams should monitor creation of any new schedu -led tasks, modifications to system services and abnormal startup scripts. In t- he cloud, continuously audit IAM roles and disable dormant accounts.

## Privilege Escalation & Credential Access: Keys to the Kingdom

What attackers or Red Teams need after persistence are high privileges or credentials. Whether that's SYSTEM privileges on an endpoint or global ad- min in Azure, the goal is escalation.

Classic techniques of privilege escalation still dominate here: dumping LSA -SS for plaintext credentials or hashes, abusing misconfigured service accoun -ts and exploiting weak Active Directory permissions. Token theft and pass- the-ticket attacks remain effective in poorly segmented networks.

Cloud attacks mirror the same principles. An attacker compromises one us -er, enumerates the environment and finds a service principal with overbroad permissions. With a stolen OAuth token, persistence and lateral movement become trivial.

**MITRE ATT&CK Mapping:** Credential Dumping (T1003), Valid Accoun -ts (T1078), Abuse Elevation Control Mechanism (T1548)

**Defensive Move:** Blue Teams should enforce least privilege, eliminate sha -red admin accounts and monitor for credential access attempts. In cloud en



ments, you should implement just-in-time access and role scoping.

## Lateral Movement: Quiet Expansion

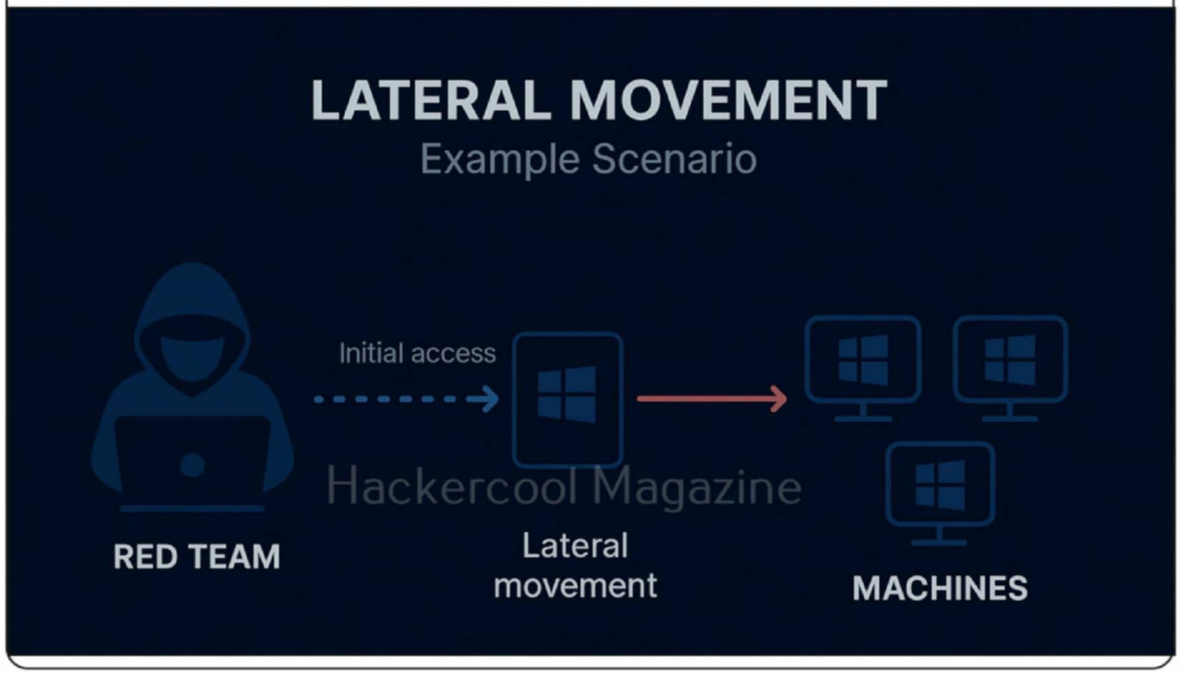

Once attackers have elevated privileges, they try to pivot to spread to the e-ntire network. Traditional lateral movement techniques like RDP, SMB, PS-Remoting, WinRM, WMI still dominate because they're built into enterprise management workflows.

The stealthiest operators don't use malware at all. They use remote manag-ement tooling already trusted by the environment like PsExec, SCCM or do-main management utilities.

In cloud and hybrid environments, attackers pivot laterally by assuming id-entities or abusing connected sync tools between on-prem and cloud identiti-es.

**MITRE ATT&CK Mapping:** Remote Services (T1021), Lateral Tool Tra-nsfer (T1570), Pass-the-Hash (T1075)

**Defensive Move:** Blue teams should segment the network and require bas-tion hosts for administrative connections. Log and alert on lateral authentica-tion patterns, particularly unusual east-west traffic or use of legacy protocols like NTLM.

## Command & Control: Blending with the Noise

The next step is Command & Control (C&C). The modern C2 landscape favors stealth over sophistication. Attackers nowadays hide in plain sight, usi-ng common protocols like HTTPS and DNS to communicate with compro-mised hosts.

Some groups leverage cloud services too (Dropbox, Slack, Telegram) as C-2 relays blending exfiltration and beacon traffic into legitimate business co-mmunications. Some other groups employ domain fronting, a technique in which malicious network traffic is disguised as benign visits to popular cloud platforms. Beacon intervals are randomized, data is chunked and payloads a-re encrypted — all to stay invisible.

**MITRE ATT&CK Mapping:** Application Layer Protocol (T1071), Web S-ervice (T1102), Domain Fronting (T1090)

**Defensive Move:** Blue Teams should implement strict egress filtering and

SSL/TLS inspection wherever possible. They should also baseline normal o-utbound patterns and alert on connections to unusual endpoints, domains w-ith high entropy or anomalous DNS tunneling.

## Exfiltration and Impact: When Data Walks Out the Door

The final stage depending on the endgame of the hacking attack is exfiltrat-ion of juicy data from the target network. Modern exfiltration is no longer c-arried as a massive data dump. It is carried in slow, quiet and disguised for-ms. Attackers are using legitimate cloud synchronization tools, compressing data with standard utilities and sending it out over encrypted channels — often via allowed APIs.

Some APTs are using steganography to hide sensitive data inside images o-r PDFs. Others exfiltrate data to compromised cloud accounts belonging to

legitimate providers, making detection even harder.

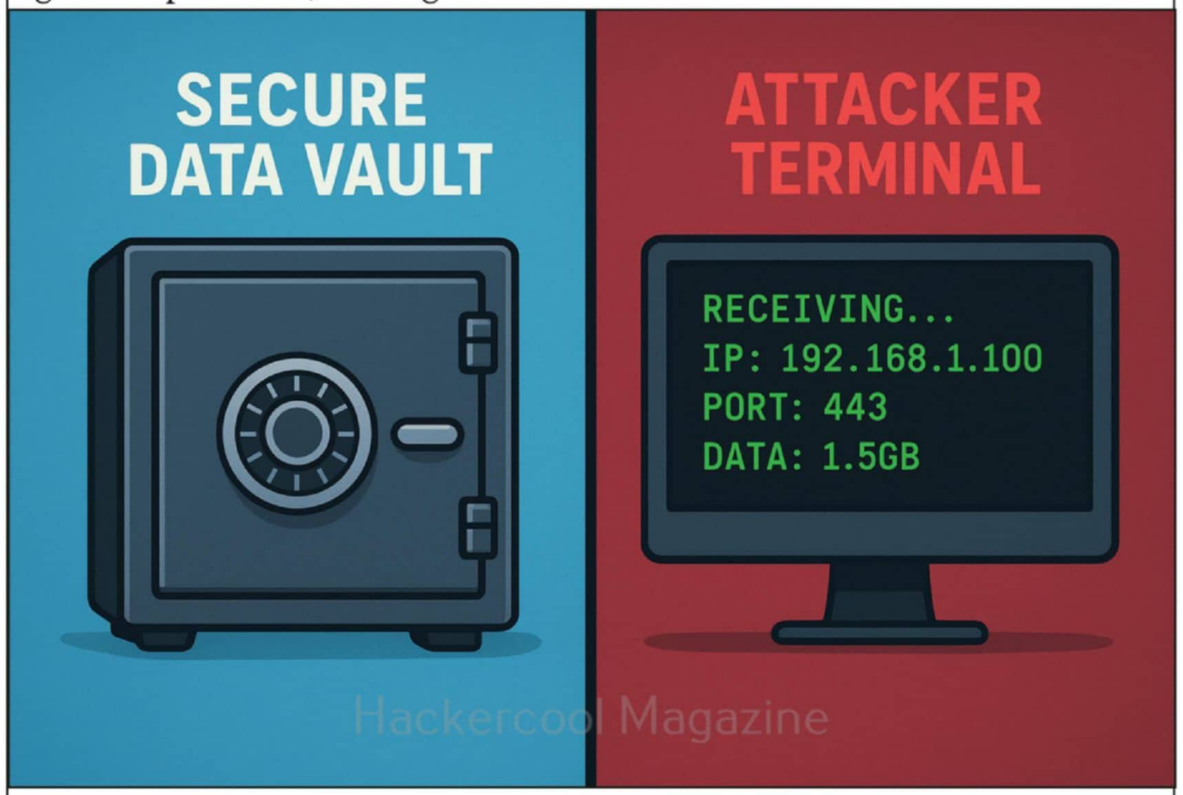


Hackercool Magazine

The endgame varies. Sometimes it's exfiltration; sometimes it's ransomware which combines encryption with exfiltration to maximize leverage.

**MITRE ATT&CK Mapping:** Exfiltration Over Web Services (T1567), Exfiltration Over C2 Channel (T1041), Data Encrypted for Impact (T1486)

**Defensive Move:** Blue Teams should use DLP with content-aware inspection and tag sensitive data for egress monitoring. Watch for irregular upload volumes, especially to new cloud endpoints. Maintain tested, offline backups — immutable, versioned and air-gapped.

## A Real-world Case Study

To see how these TTPs can be combined together for a successful hack, let's see an example of a recent red team exercise of an organization named LOOKRECKAH (name changed).

Hackercool Magazine

**Reconnaissance:** The Red team while performing reconnaissance found that the company's job listings mentioned "Active Directory Federation Services" and "Citrix Gateway," revealing parts of its infrastructure.

**Initial Access:** To gain initial access, a spear phishing email was sent to targeted IT staff with a fake MFA enrollment prompt, capturing valid credentials via an AiTM proxy.

**Execution:** The Red Team used these credentials to log in through the VPN, executed some PowerShell commands to query domain controllers and loaded an in-memory C&C agent via rundll32.

**Persistence:** For persistence, they created a hidden scheduled task to re-establish access every 24 hours.

**Privilege Escalation:** They then performed POST-exploitation enumeration. While doing this, a misconfigured service account with domain admin privileges was discovered via BloodHound.

**Lateral Movement:** Using that account, they deployed a payload via PsExec to internal file servers.

**C2 & Exfiltration:** Data was exfiltrated over HTTPS to a Dropbox accou
-nt – all within allowed egress policies of the organization.

If you noticed the above real-world example keenly, you can see that the a
-ttackers did not exploit any zero-day vulnerabilities. Each and every step rel
-ied on configuration oversights and predictable human behavior.

## Why These TTPs Still Work

You know why these TTPs still work in real-world now. The truth is a bit u
-ncomfortable for defenders. They already know these techniques exist but c
-an't prevent them without breaking workflows or overwhelming analysts. At
-tackers thrive in that gray zone – using tools that are "too normal to block"
and "too quiet to alert."

Patch cycles get lagged, users click links, credentials get reused and visibility stops at network boundaries that no longer exist. The modern kill chain endures because enterprise complexity creates endless opportunity for



adversaries who are patient, methodical and adaptive.

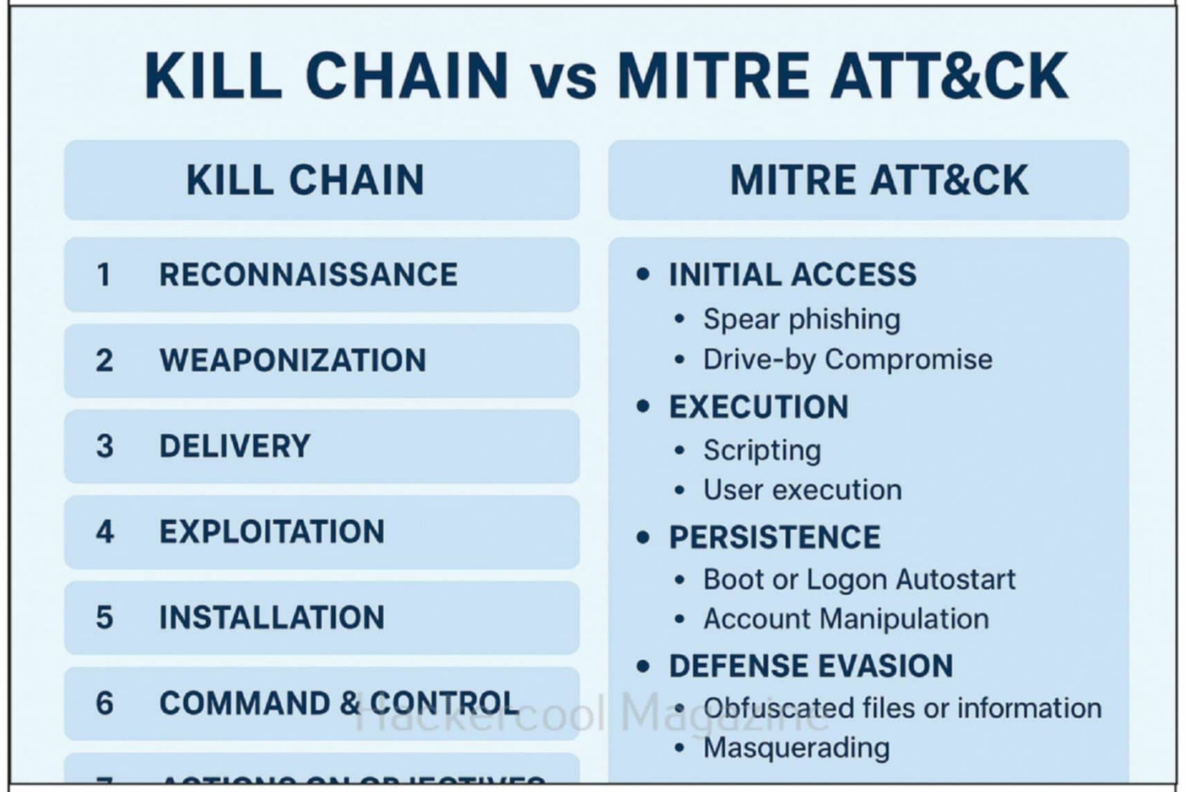


## How to Build a network resilient to Modern TTPs

To break this kill chain, Blue Teams must accept that focusing on preventi-
on alone is a losing game. The goal isn't just perfection – it's early detection
and rapid disruption. Here are 5 steps you can take that help in building mo
-dern resilience.

**1. Instrument Everything:** Centralize logs of endpoints, identity, and clo-
ud. Telemetry is of no use if it's siloed.

**2. Baseline Behavior:** Make a note of what "normal behavior" looks like
in a network like logon patterns, outbound traffic, admin tool usage and flag
any deviations from it.

**3. Embrace Purple Teaming:** Regularly simulate APT-style operations m
-apped to MITRE ATT&CK to test controls.

**4. Limiting Lateral Movement opportunities:** Enforce least privilege, s-egment networks and implement just-in-time admin access.

**5. Plan for the Inevitable:** Despite taking all defensive measures, assume there will be a compromise or breach. Practice containment and recovery as rigorously as prevention.

## DEFENSIVE DETECTION POINTS

Overlay of detection controls mapped to each phase of the kill chain

| | Reconnaissance | Weaponization | Exploitation | Actions on Objective |
|---|---|---|---|---|
| SIEM | ✓ | ✓ | ✓ | ✓ |
| EDR | ✓ | ✓ | ✓ | ✓ |
| IDS/IPS | ✓ | ✓ | ✓ | ✓ |
| Antivirus | ✓ | ✓ | ✓ | ✓ |
| Network Traffic Analysis | ✓ | ✓ | ✓ | |

Hackercool Magazine

The organizations that recover fastest aren't the ones with the biggest budg-ets — they're the ones that expect failure and build muscle memory around i-t.

## Conclusion

The modern kill chain isn't dead yet, it's a reminder. It's a reminder that at-tackers still rely on the same fundamental phases because human systems, o-perational shortcuts etc haven't evolved as fast as defenses.

"Blue Teams don't chase alerts; they chase the truth hidden inside them."



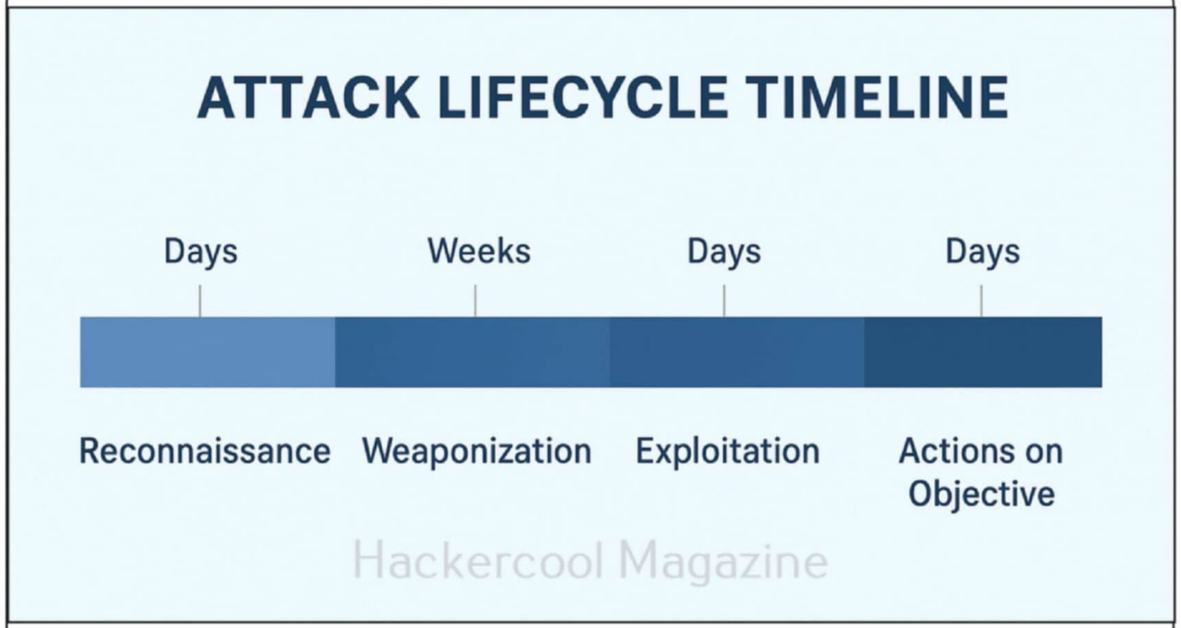


Red teams exploit that reality. Blue teams must also embrace it. The most successful security programs aren't trying to eliminate risk entirely; they're b-uilding the visibility, agility and trust to respond faster than adversaries can a-dapt. Because in the end, the kill chain isn't a model — it's a mirror. It show-s us where we're weakest and where we can get stronger.

*This part of the page has been intentionally left blank*

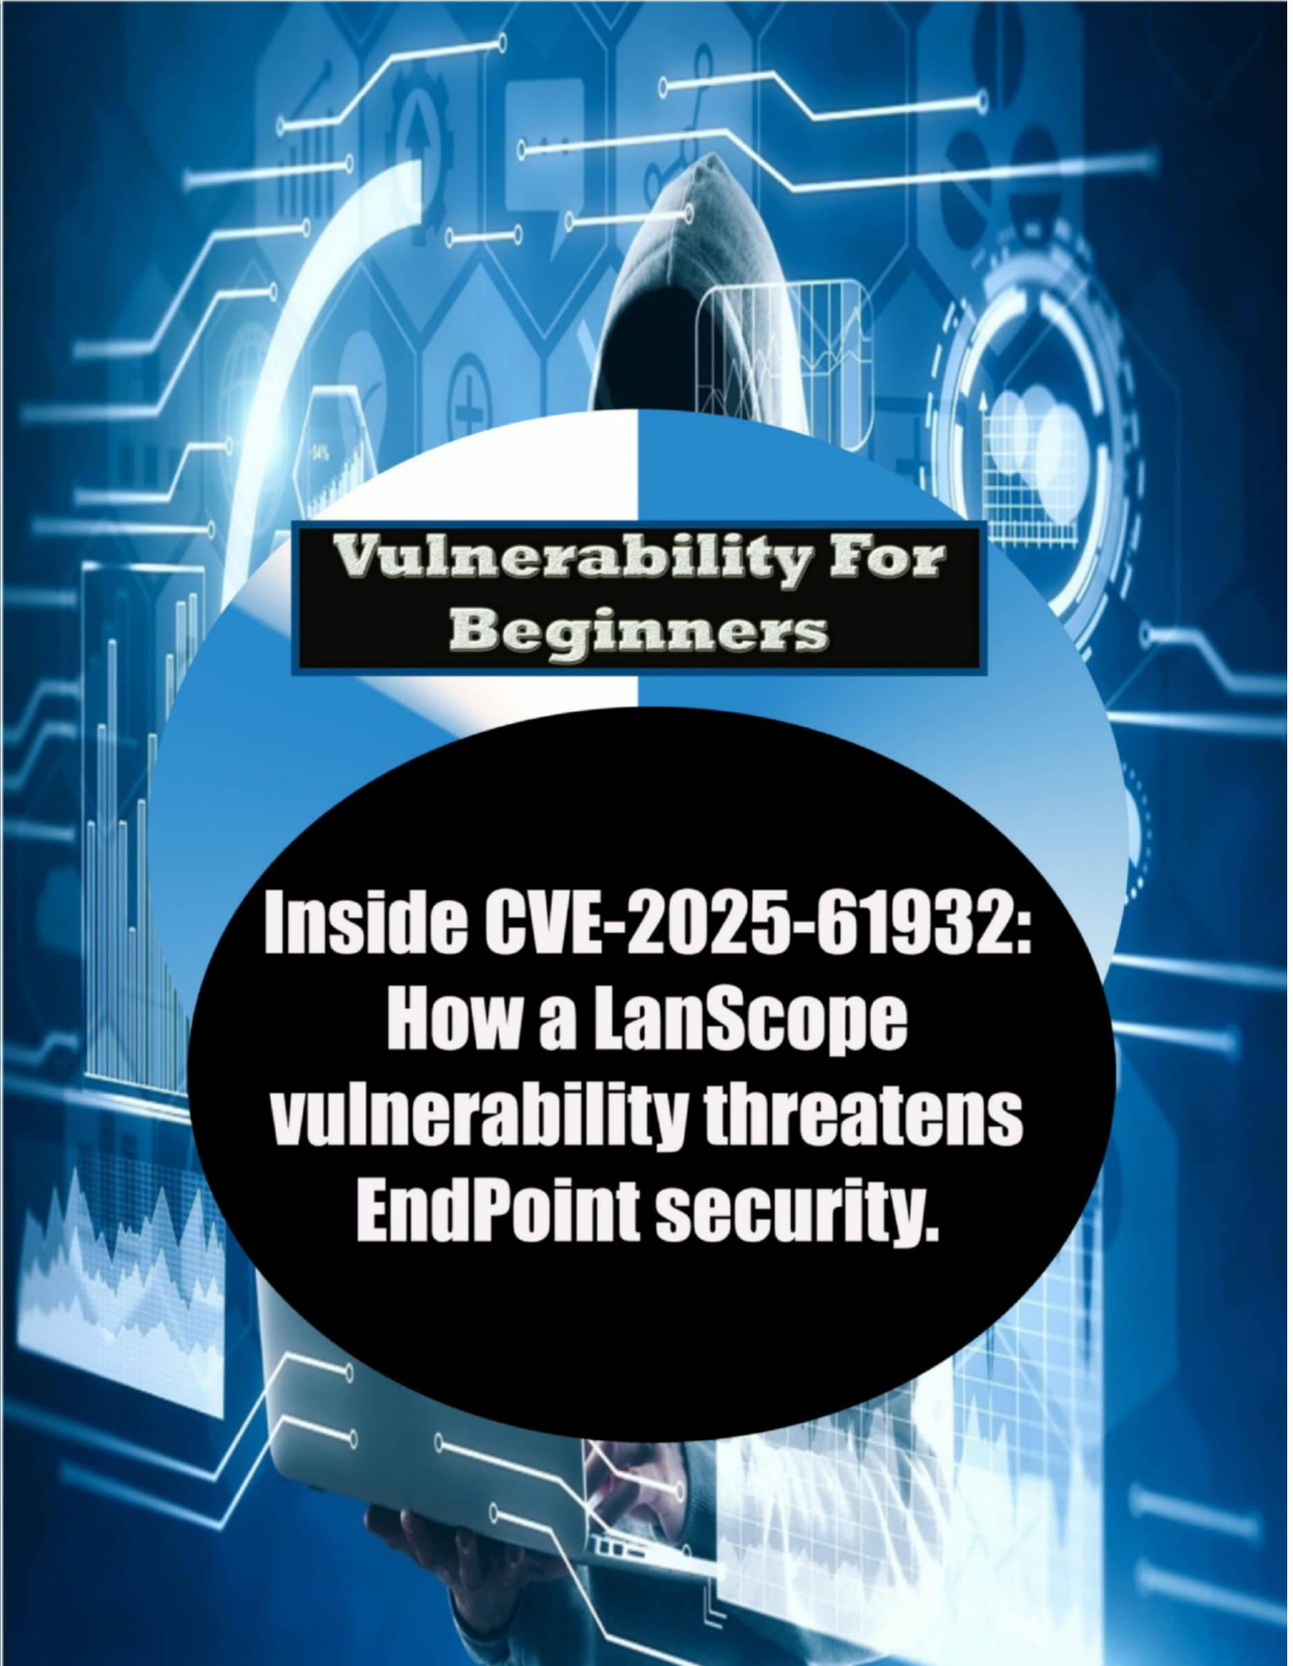




A severe vulnerability has been disclosed by Motex in its LanScope Endpoint Manager that is already under active exploitation in real-world.

Motex LANSCOPE Endpoint Manager

Hackercool Magazine

**Remote code execution vulnerability**

## About the vulnerability

The vulnerability being tracked as CVE-2025-61932, is a remote code execution vulnerability that affects On-Premise edition of LANSCOPE Endpoint Manager. It has CVSS 3.0 rating of 9.8.

The vulnerability does not affect the centralized management console but is present in the code of its Client Program (PR) and Detection Agent (DA) which are installed on endpoint devices to collect data to be sent to centralized management console.

The flaw is due to an improper verification or validation of the source of communication channel flow. In simple terms, it means this vulnerability is a result of performing inadequate checks on incoming communication packets.

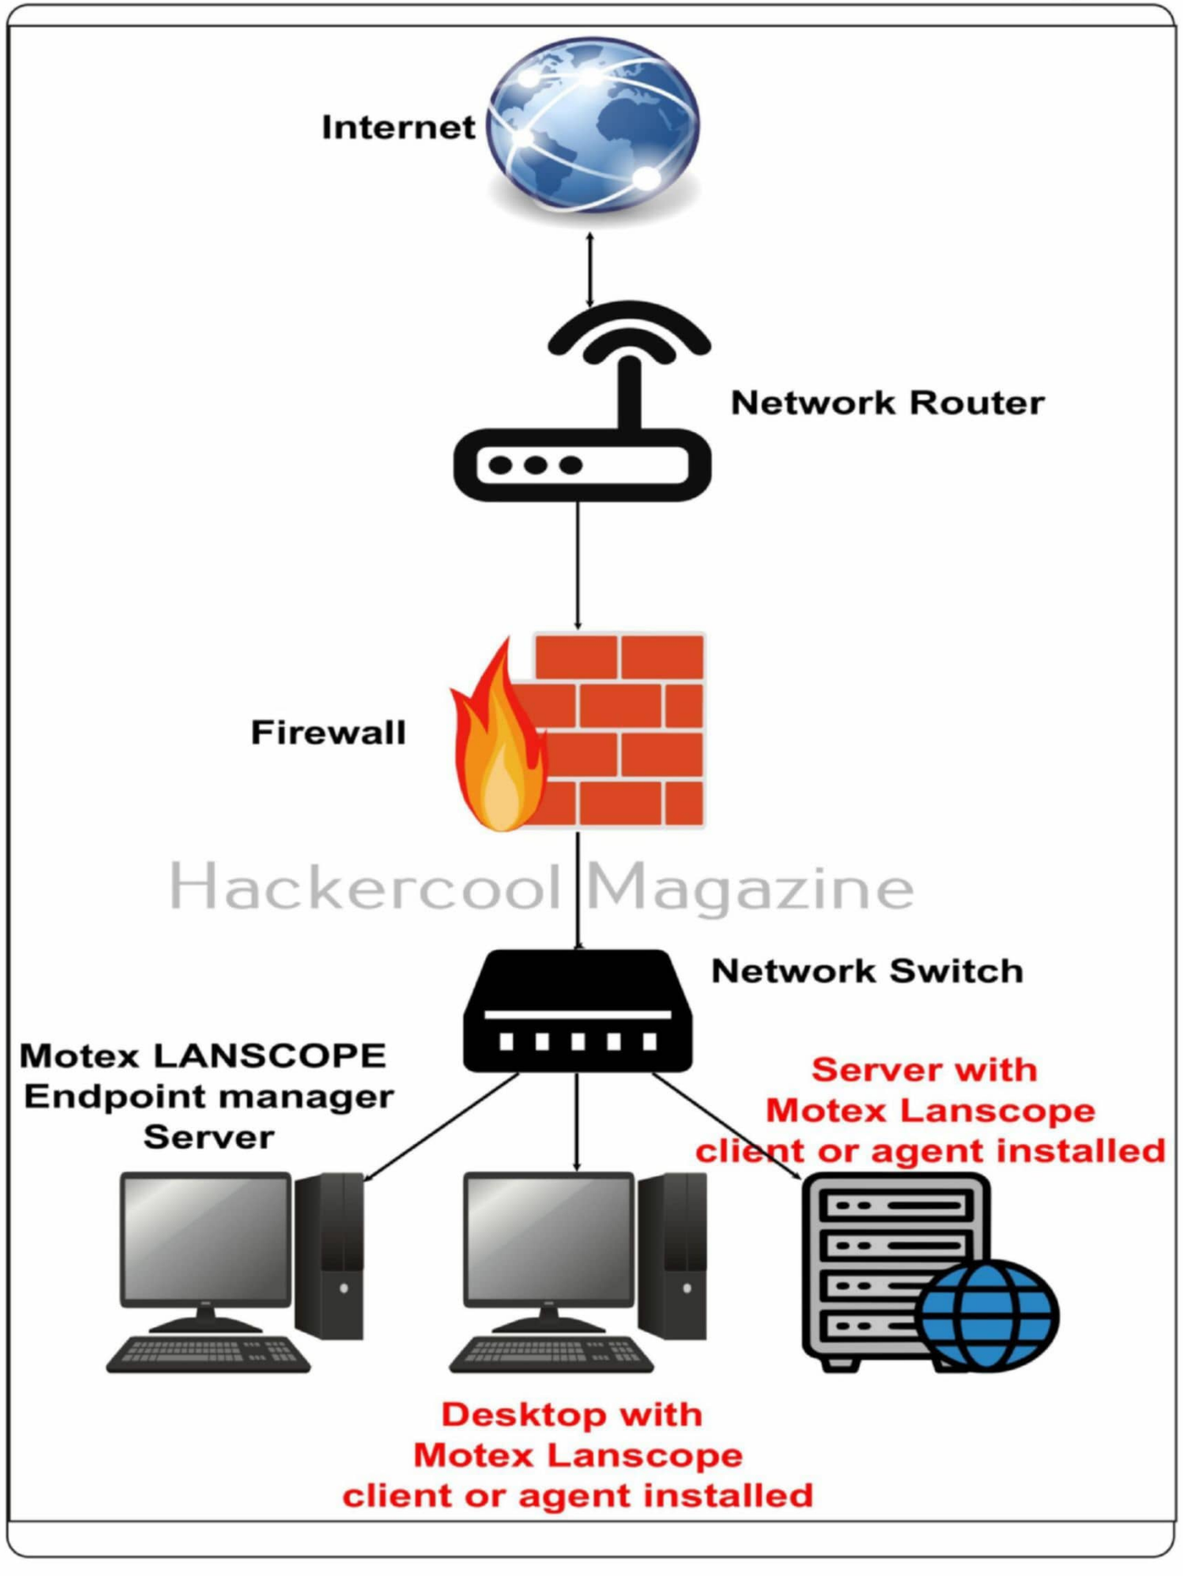




LANSCOPE On-Premise Edition with versions up to 9.4.7.1 are vulnerable. LANSCOPE also has a cloud-based edition that is safe from this vulnerability.

## How can hackers exploit this vulnerability?

This vulnerability can be exploited remotely by attackers by sending a specially crafted packet to the exposed instance of vulnerable device that allows them to bypass authentication and gain access to the device.

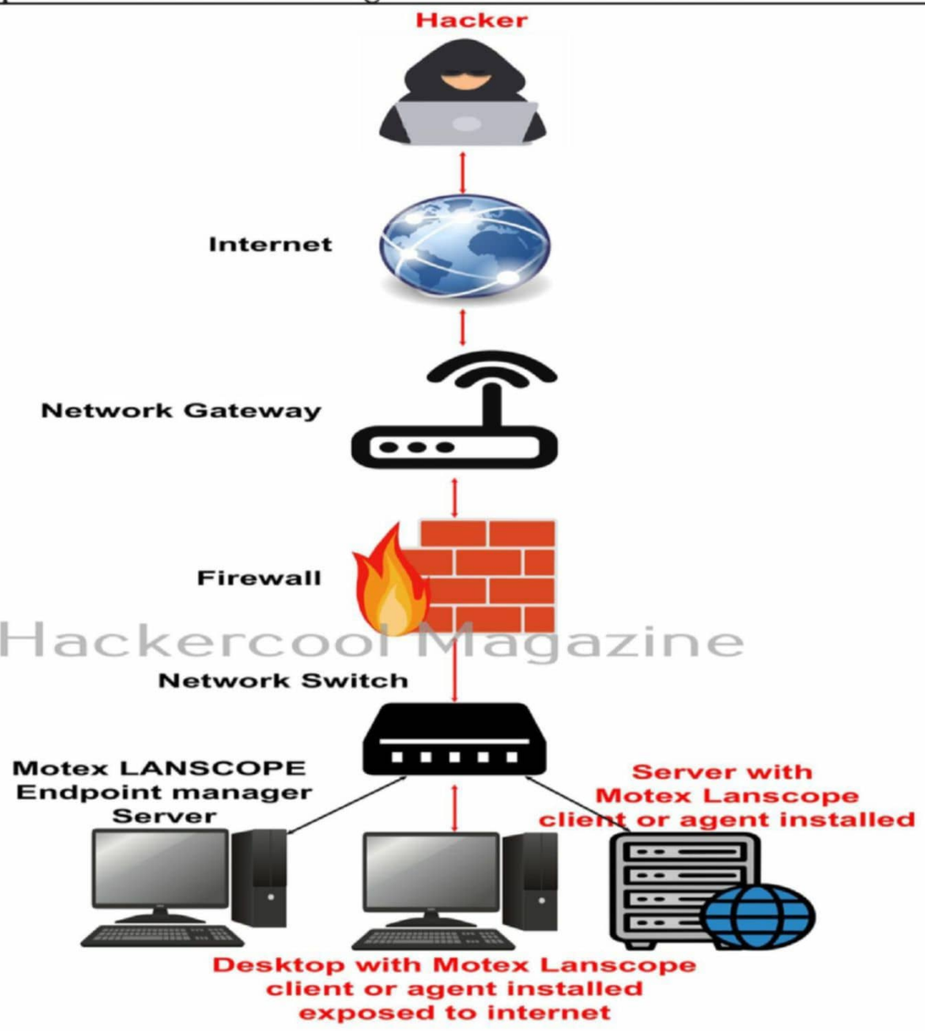


Fig: Exploitation scenario of CVE-2025-61932 vulnerability in Motex LANSCOPE Endpoint manager

All attackers have to do is impersonate legitimate sources while sending this crafted packet.

## Impact

LANSCOPE EndPoint Manager is a very popular tool used for managing I-T assets across networks. They are mostly used in financial, healthcare enter-prises. Like any other remote code execution vulnerability, it allows attacke-rs to execute malicious code on affected systems remotely and fully compro-mise the endpoint device.

Since the vulnerability affects the client software of LANSCOPE Endpoint Manager which can be installed on any number of devices in an organizatio-n, the exploitation potential is very high if any of these devices is exposed to the internet. Add to this the point that to exploit this vulnerability attackers don't need any user interaction at all.

Once attackers exploit this vulnerability, they can pivot to other devices, d-eploy malware, ransomware or steal sensitive data.

Vulnerabilities in endpoint management tools are juicy fruits for attackers a-s they often run with elevated privileges and offer an easy pivoting option. CISA has issued an immediate alert and has already added this vulnerability to its Known Exploit & Vulnerabilities (KEV) catalogue since it is being acti-vely exploited in the wild.

It is estimated that active exploitation may be occurring since April 2025. S-ome real-world exploitation cases including this vulnerability are Qilin rans-omware attack on Asaki brewery and Askul e-commerce retailer.

## How to manage this vulnerability?

Motex has already released a fix for this vulnerability and users of this prod-uct should immediately apply this fix as soon as possible. You can apply this fix from LAN SCOPE PORTAL, which is their customer support portal.

Note that this update is for client PC's only. The central management conso



-le does'nt need any fix. The fixed versions of the software are 9.3.2.7, 9.3.3.9, 9.4.0.5, 9.4.1.5, 9.4.2.6, 9.4.3.8, 9.4.4.6, 9.4.5.4, 9.4.6.3 and 9.4.7.3.

Make sure your version of LANSCOPE client is any one of the above. Also note that this vulnerability has no workaround and only way to fix it is to ap-ply vendor supplied patches as soon as possible.

*This part of the page has been intentionally left blank*

Security Onion has been a cornerstone of open-source detection and respo
-nse for years. With the 2.4.180 release, the project stops short of reinvention and instead sharpens the tools SOCs and labs already rely on. This release will give you better analyst ergonomics, improved host visibility, scalable me
-ssage pipelines etc. This update helps not only SOC engineers who need re
-liability at scale but also to teams building red team and purple team labs w
-ho want realistic, enterprise grade telemetry without the shock of enterprise sticker.

Many SOC upgrades are measured in micro improvements: a tweak to a UI here, a library bump there. 2.4.180 bundles a set of focused, practical ch-anges that reduce friction for analysts and platform owners. Whether you op
-erate in AWS/Azure, on prem, or inside nested lab VMs, the release addres
-ses three oft complained pain points:

- onboarding and learnability (contextual help)
- visibility gaps (offline agent alerts, static hostname mapping)
- scale and integration (Kafka outputs, component upgrades, hypervisor ma

-nager features)

These changes help close the gap between a community SOC deployment and a hardened, repeatable enterprise stack.

## Latest features to be excited about

## 1) Contextual Help — UX that actually helps

The interface of Security Onion now includes on screen contextual help bu -ttons tied to major views. Analysts, especially novices, if using this release w -ill spend less time reading and trying to understand documentation and mor -e time triaging. For training programs and tabletop exercises, this reduces fri -ction and keeps the focus on detection skills rather than UI guesswork.

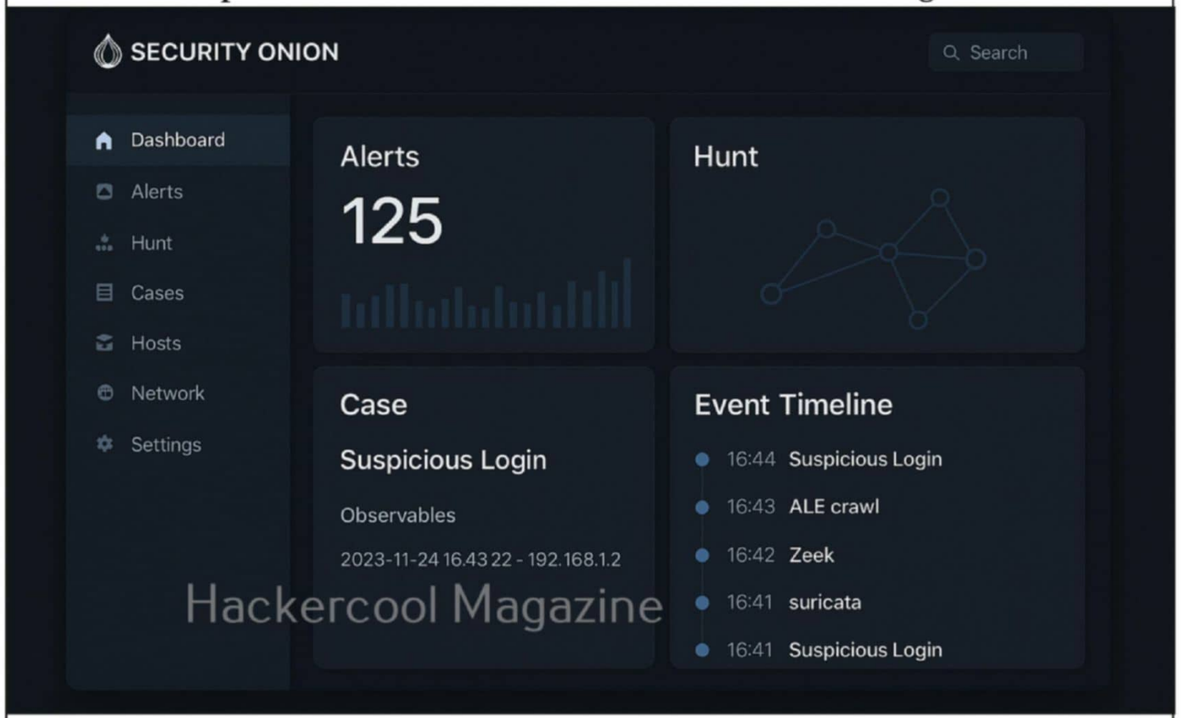


## 2) Offline Agent Alerts — detect the silence

Just imagine a scenario where you are monitoring the security of the netwo



-rk and one of your agents on client computers goes offline. Don't worry. Security Onion now shows an alert when expected agents goes offline. This is a fundamental requirement: telemetry that suddenly stops is often the first sign of tampering, agent removal or network issues.

Just imagine if an attacker disables the endpoint agent as part of lateral mo- vement. Offline alerts can kick off an early investigation rather than relying solely on noisy detections after the fact.

## 3) Cancel Long Running Queries — take back control

Analysts can now cancel their own long queries (alerts, hunts, dashboards) with an on screen control. This improves responsiveness and reduces resour- ce contention on heavy deployments.

## 4) Static Hostname Mapping (no RDNS required)

Many lab and segmented networks lack reverse DNS. Manual/static hostna -me mapping allows administrators to present human readable hostnames in alerts and dashboards without requiring DNS infrastructure.

## 5) Kafka Output Policy & Component Upgrades

In the latest release, a Kafka output policy lets Security Onion stream alert/ events into scalable messaging pipelines. This is ideal for high volume, multi tenant environments. The release also includes updated core components fo -r performance and security. This can help in easier integration into cloud in -gestion pipelines, SIEM aggregation stacks or downstream analytics.

## 6) Hypervisor / Manager Enhancements (Pro)

In the PRO version, Manager nodes with improved hypervisor capabilities allow consolidation and nested VM hosting of sensor components. This is es -pecially useful for lab builders and enterprises wanting to emulate distribute -d deployments on limited hardware.

## 7) Reporting & Exporting (Pro)

Pro edition also adds built in reporting and data export features–PDF and custom exports are now part of the platform's offering. That's a direct win fo -r incident review, executive dashboards, and compliance reporting.

## Other important updates

## Updated Tools

Most of the analysis tools have been update to the latest versions. These include Suricata, Zeek, Elasticsearch, Kafka etc.

If you crave reduced analyst friction and stronger host agent awareness, this release of Security Onion pays rich dividends. The Pro features are attractive for enterprises, but even the community edition gives you tangible benefits.

A critical vulnerability has been discovered in Ubiquiti Uni-Fi Access applications which can be exploited by attackers without the need of any authentication.

## About the vulnerability

The vulnerability tracked as CVE-2025-52665 affects the Unifi Access application versions from 3.3.22 to 3.4.31. The vulnerability is due to a misconfiguration introduced in version 3.3.22 and has been assigned a CVSS V3.1 score of 10.0.

The vulnerability exists in /api/core/backup/export endpoint. This endpoint handles backup operations. However, this endpoint accepts a directory parameter (dir) without any sanitization. Moreover, the backup operation should only be listening on IP 127.00.1 (local host with restricted access) but it is exposed on port 7780). This misconfiguration allows attackers to access it with-



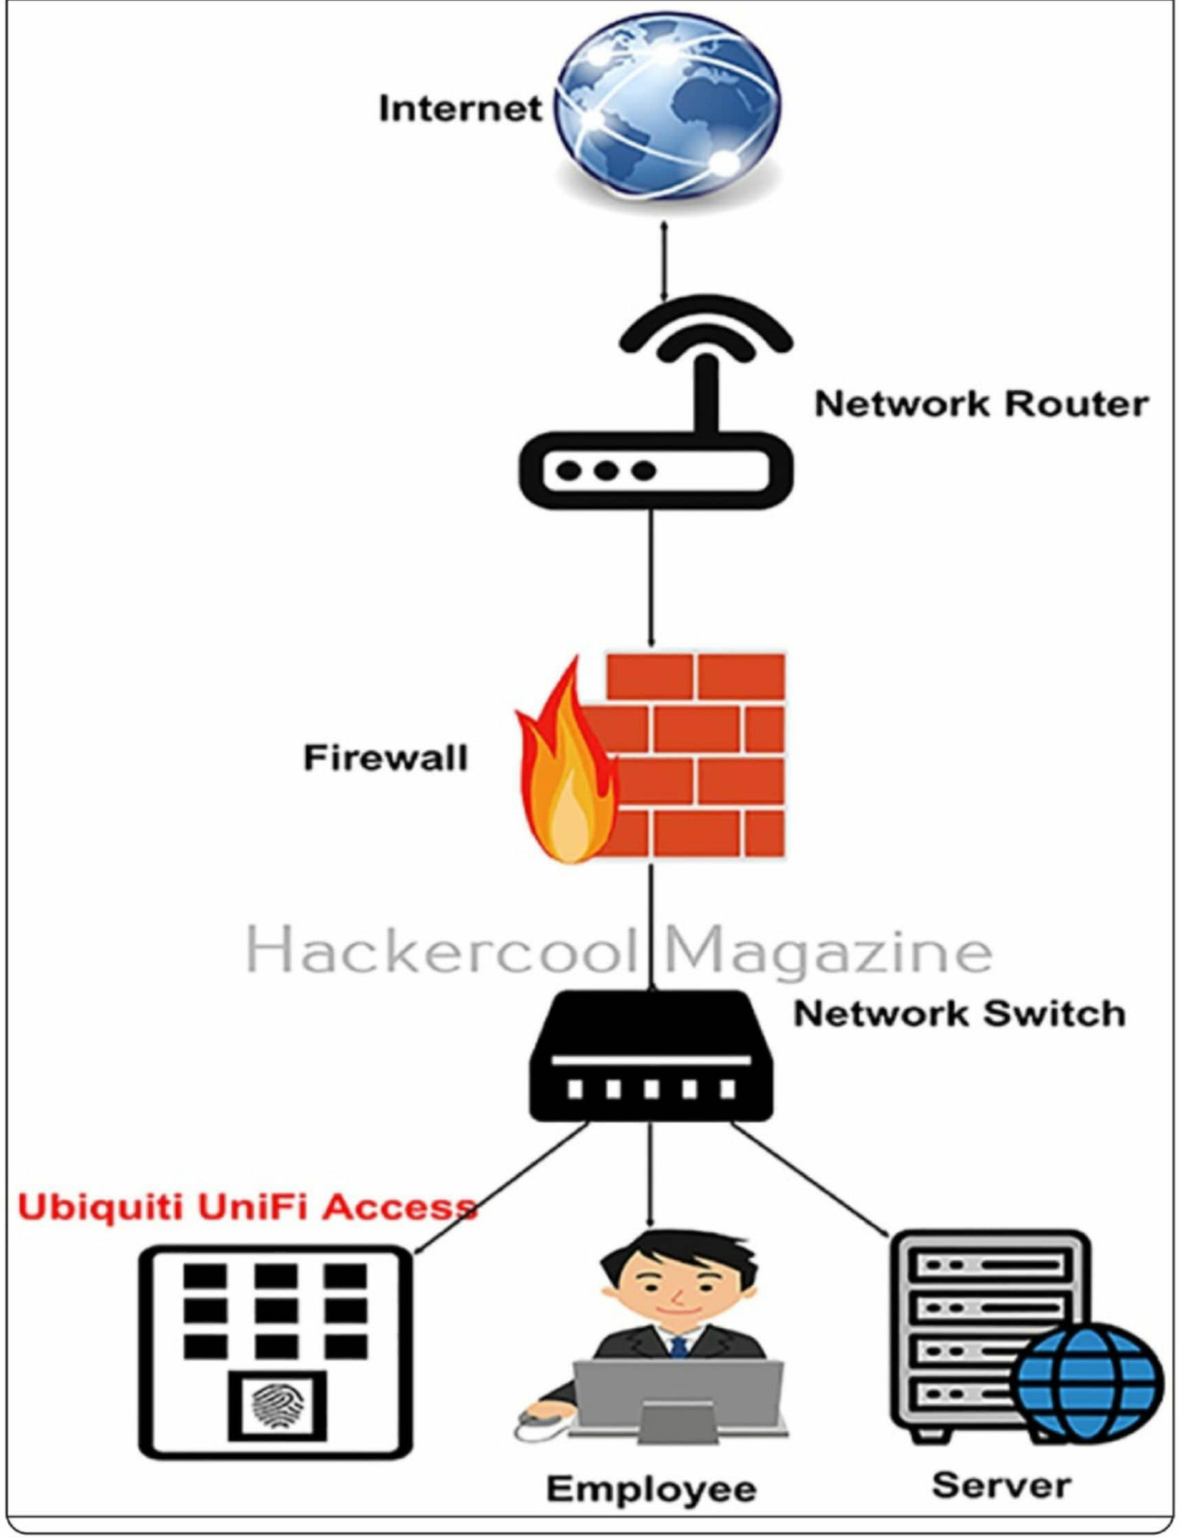

out any authentication.

## How hackers can exploit this vulnerability?

Attackers can exploit this vulnerability remotely if this device is exposed to internet either due to a misconfiguration or with intention. It can also be exp-loited if hackers already have network access or if they have breached peri-meter defense.

Exploitation begins by attackers scanning for exposed UniFi devices on its default ports. Once they find it, they can exploit this by injecting shell metac-haracters into the 'dir' parameter. Then they escape the command shell and execute RCE commands.

Attackers send a malicious formatted JSON POST request to the target en-dpoint imitating legitimate backup orchestration calls to exploit this vulnera-bility.

## Impact

By exploiting this vulnerability, attackers can fully take control of door acc-ess systems and they can change access controls, unlock doors etc thus affec-ting physical security of the organization.

Exploitation of this vulnerability doesn't require any privileges and can esp-ecially be dangerous in connected setups like smart buildings and offices th-at rely on IoT devices to ensure physical security.

## How to manage this vulnerability?

Ubiquiti has already released a patched version to mitigate this vulnerabilit-y. Users are advised to update this software ASAP to prevent its exploitatio-n. The version of UniFi OS safe from this vulnerability is 4.0.21 or greater.

There should also be regular network auditing of network configurations to see if any UniFi devices are exposed to internet. Moreover, IoT and access control software should be protected with strong authentication.



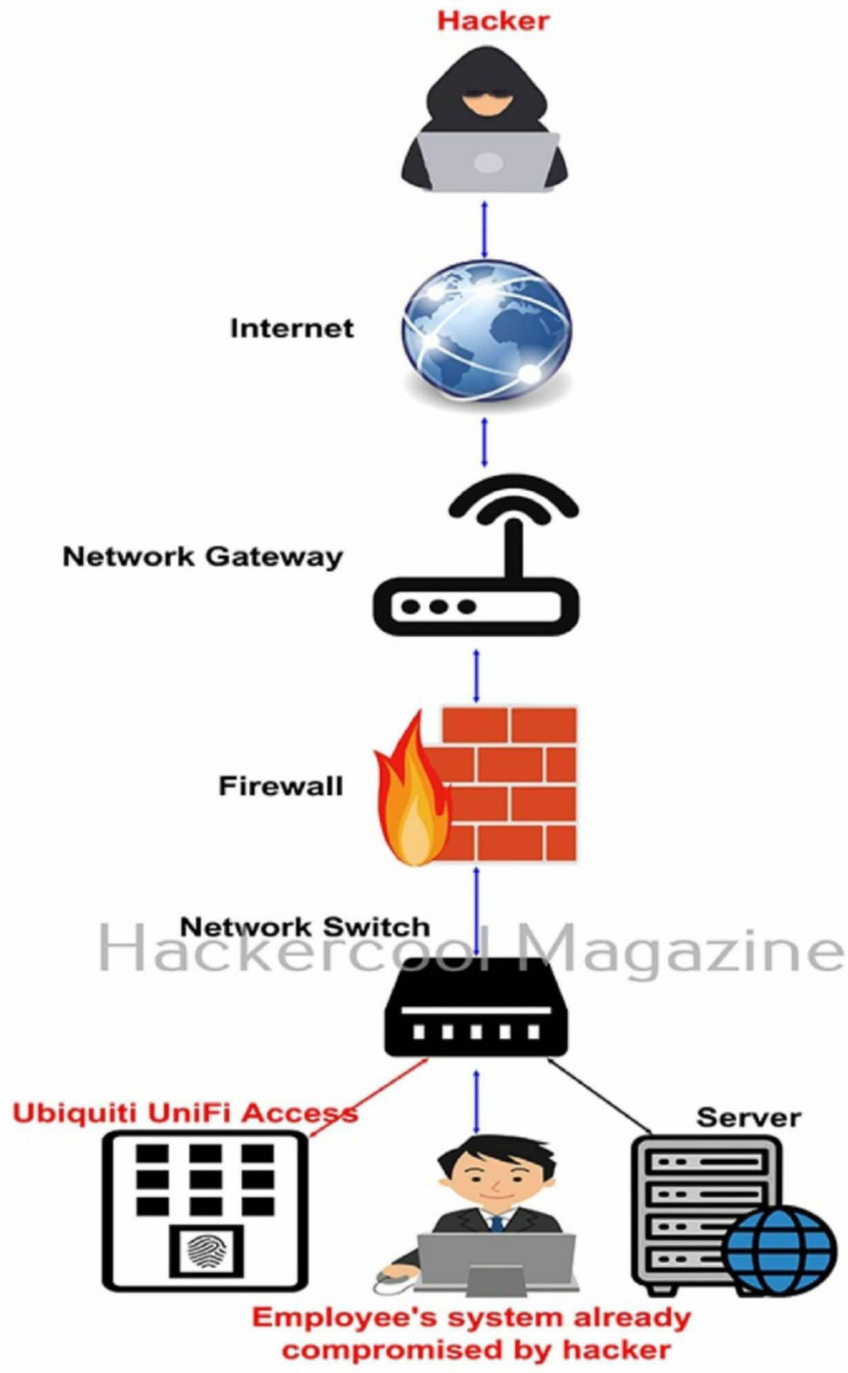


Fig: Exploitation scenario of CVE-2025-52665 vulnerability in Ubiquiti's UniFi OS

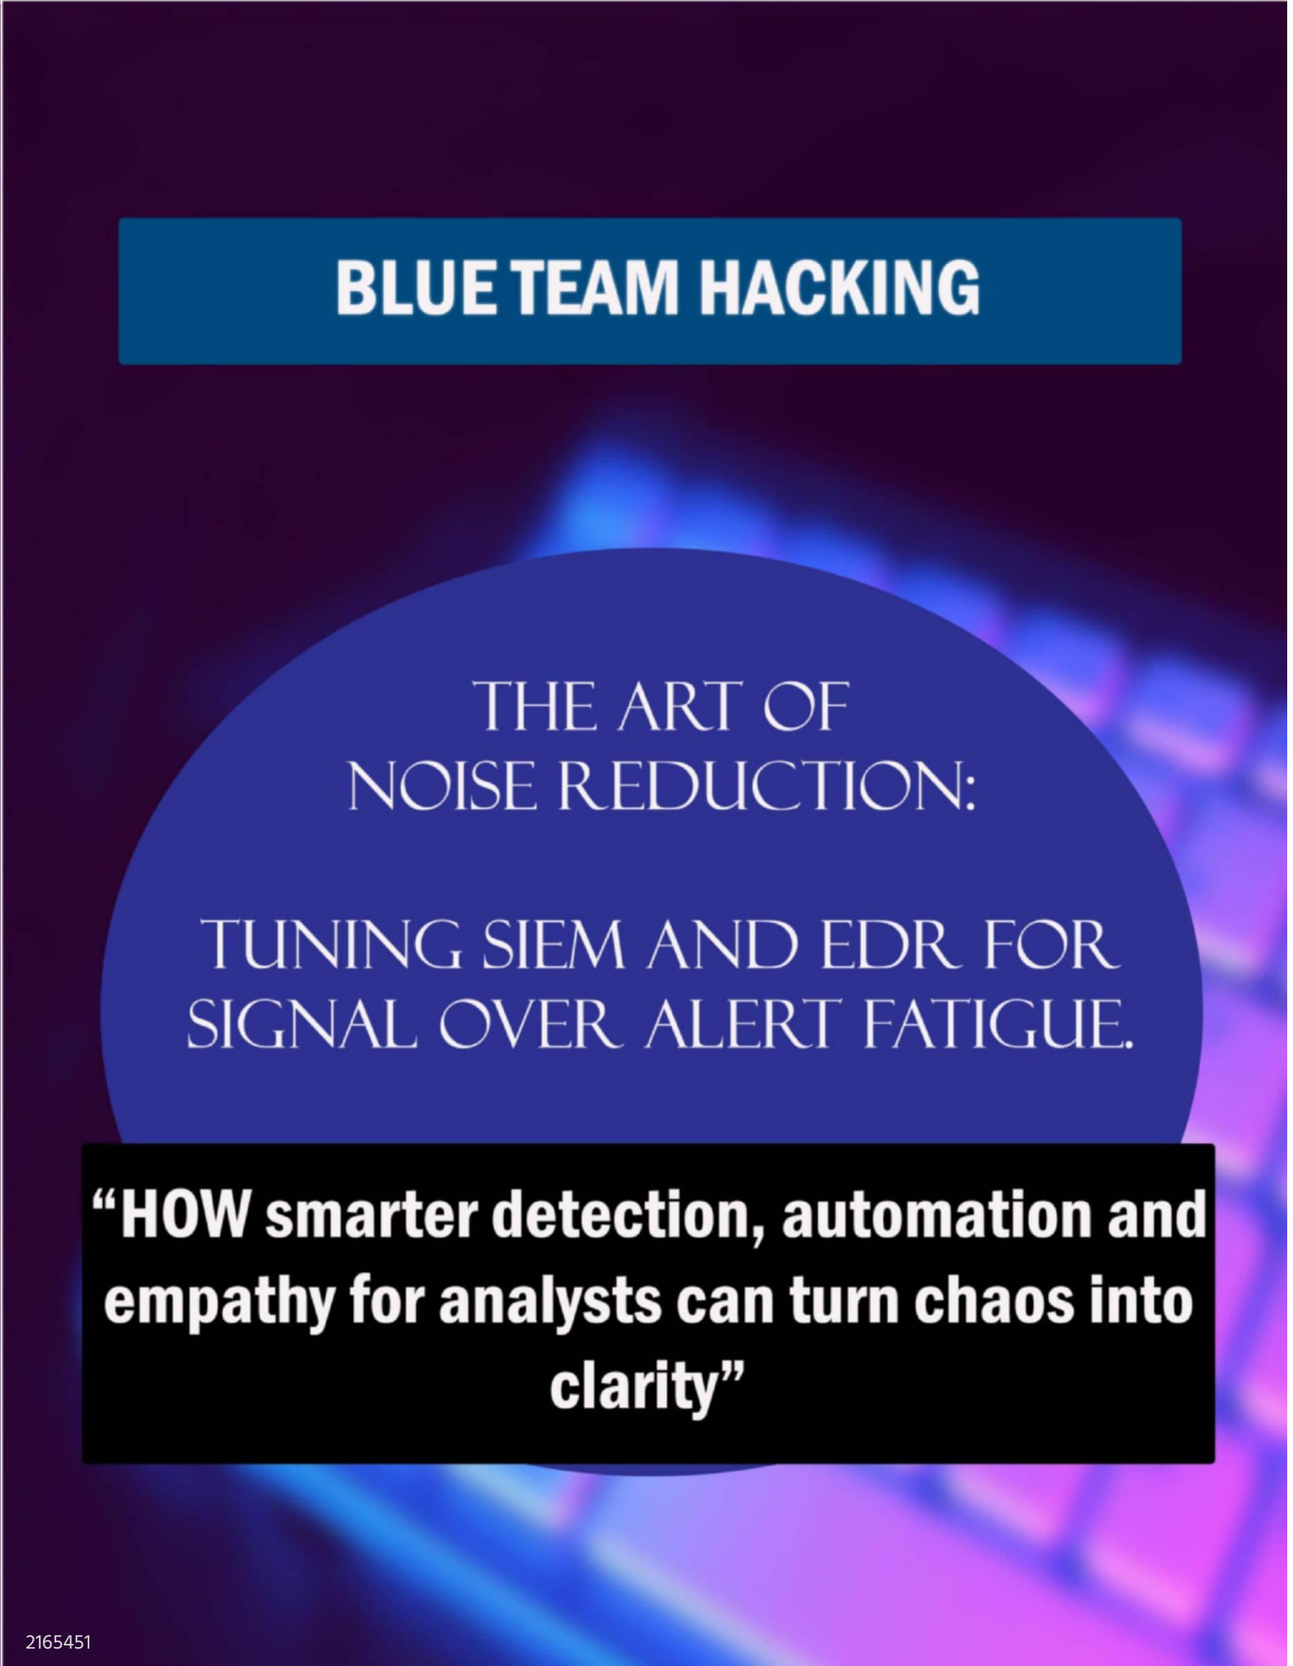




## The Problem No One Wants to Admit

What according to Blue Teams is the best way to defend the organization? Set up SIEM's, EDRs, firewalls and other defenses and have a central management console to view the alerts and logs they generate. Any malicious activity can be detected by noticing the alerts and action taken upon. How idealistic scenario? However, like Ohm's law in physics, some things look better in ideal conditions. Practically, it is different. It is true in the case of Blue teaming too.

Every SOC analyst knows the sound — the endless chime of alerts, the flicker of dashboards lighting up like a Christmas tree. Every day, thousands of events stream in from SIEMs, EDRs, firewalls, identity platforms and cloud services. The average enterprise logs tens of billions of security events per month. Out of these tens of billions, maybe a few dozen alerts matter to the se- curity of the organization.

Yes, that is a gross imbalance. But it's reality. This imbalance has a name. It's called alert fatigue.

Analysts are getting drowned in low-value noise which includes misfired ru-les, duplicate alerts or events that lack context. Many teams spend 80% of th-eir day triaging false positives rather than hunting genuine threats. It's not ju-st inefficient, it's dangerous too. When everything is "critical," nothing actua-lly is.

In BLUE TEAM HACKING feature of this month's Issue, we bring you th-e Art of Noise Reduction i.e how Blue Teams can reduce unnecessary nois-e (i.e unnecessary alerts and logs) and improve their detection stack so that t-hey can focus on what truly matters.

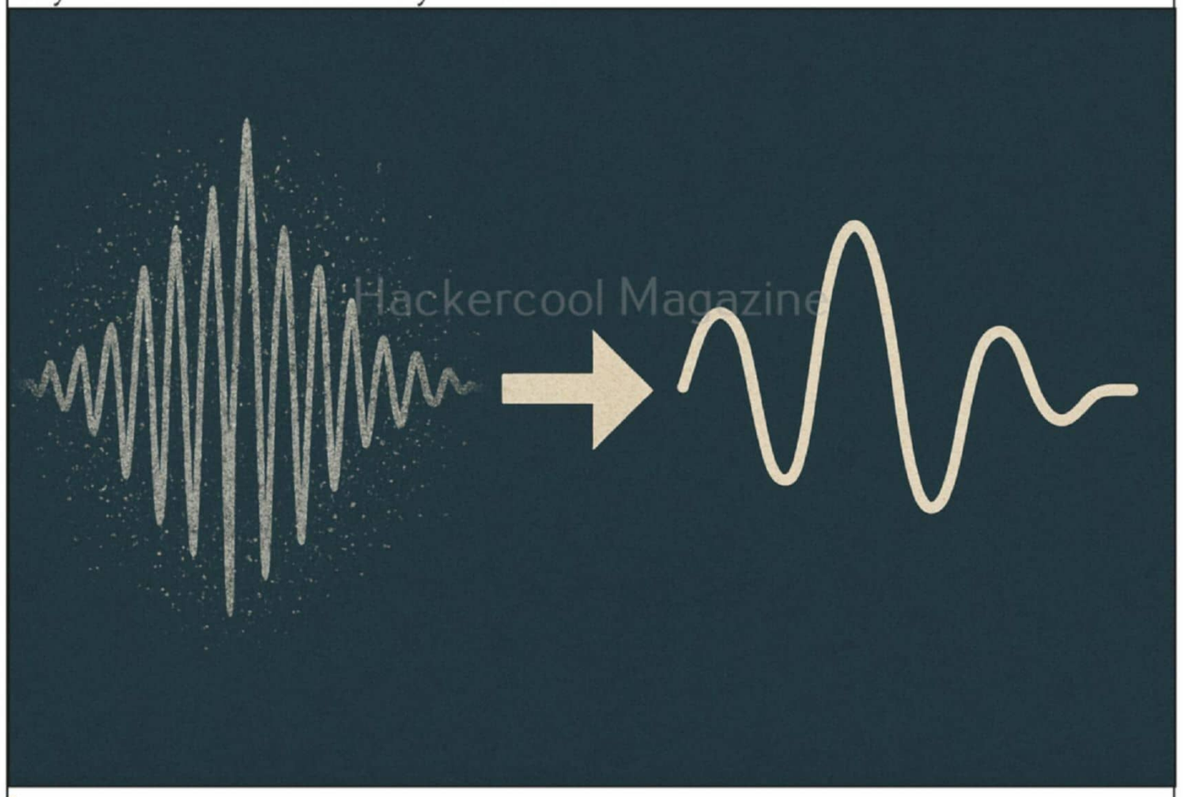


## Why Alert Fatigue Happens?

If you ask me, alert fatigue is the natural byproduct of good intentions gon-e unchecked. You will soon see why. Some of the causes of alert fatigue are,



## 1. Too Many Tools, Too Little Integration

Modern SOCs are built from a patchwork of technologies: SIEMs, EDR/X-DR platforms, cloud-native analytics, identity protection, network sensors an-d custom detections. Each tool comes from different makers, with its own lo-gic and thresholds. Few talk to each other gracefully.

Without normalization and correlation, every tool shouts out independentl-y — and the human in the loop is left to make sense of it all.

## 2. Default Rules, Default Pain

All the above products get shipped with their own default rules. They may be in 100s or even 1000s. They're designed with good intention to be a starti-ng point for new users. Many Blue teams enable them all "just to be safe." However. these generic rules rarely match the organization's requirement an-d behavior, triggering false positives by the thousands.

## 3. Lack of Context

A brute-force alert from an internal server looks terrifying until you realize it's a vulnerability scanner doing its job. SIEMs often flag behavior in isolati-on, without linking related activity or enrichment from threat intelligence.

## 4. Human Burnout

In some cases, overworked analysts start getting tired, either mentally or lit-erally. Fatigue breeds complacency and critical signals slip through. The hu-man cost of constant triage is high: burnout, turnover and erosion of trust in detection systems.

## Refocus the Goal: From Volume to Value

By now, you have understood the various reasons for alert fatigue. Now, le

-t's learn what you can do about reducing alert fatigue. Noise reduction and thus reducing alert fatigue isn't all about muting alerts; it's about raising fidel -ity. Instead of counting how many alerts they are receiving, SOCs should fo -cus on measuring how many alerts lead to verified action.

High-fidelity detection = meaningful + actionable + validated.

That means building detections that are aligned to real adversary behavior (MITRE ATT&CK is a great baseline) and using feedback loops to refine th -em continuously is important to reduce noise. Think of it as detection engin -eering, not just configuration.

Hackercool Magazine

**Fig: Steps to reduce alert fatigue**



To reduce alerts and at the same time improve value of the alerts, you nee-d to take some steps. They are,

1. Know your data (and cut the bloat)
2. Tune Your SIEM for Context and Correlation
3. Tune Your EDR/XDR for Precision
4. Automate Intelligently
5. Measure What Matters

Let's learn about each of these steps in detail.

## Step 1: Know Your Data (and Cut the Bloat)

Before trying to tune any alerts, you must first understand the data feeding them. For this you need to inventory your data sources. Make a list of every log source in your SIEM or EDR. For each, answer:

- What value does this data bring?
- How often does it trigger alerts?
- Is it redundant with other sources?

You'll often discover unused telemetry eating storage and budget. If you're paying to ingest logs no one ever queries, turn them off or downsample the-m.

**Normalize and Enrich:**

Use a consistent schema (such as Elastic Common Schema or Open Teleme -try) to align fields across platforms. Enrich raw logs with context – hostnam -e, user identity, location and business criticality. A simple enrichment step c -an raise alert fidelity dramatically.

**Apply Data Quality Metrics:**

If timestamps drift or fields are missing, your correlation rules will crumble. So, monitor ingestion health as diligently as you monitor detections themsel-ves.

## Step 2: Tune Your SIEM for Context and Correlation

The SIEM should be a signal amplifier, not a megaphone. Proper tuning tur

-ns disjointed events into coherent stories.

**Build Tiered Rules:**

Separate detections into tiers as shown below.

- Tier 1: Noisy but valuable (e.g., failed logins, port scans). Send to dashbo -ards or daily reports.
- Tier 2: Medium-confidence detections needing enrichment. Send to triag -e queue.
- Tier 3: High-fidelity correlations that likely indicate compromise. Escalate immediately.

This triage design prevents low-level events from clogging the pipeline whil -e keeping visibility.

**Correlate with Purpose:**

Use correlation rules to connect related events — for example:

- Multiple failed logins followed by a successful login from a new geogra- phy
- PowerShell script execution network connection credential dump

Context transforms noise into insight. MITRE ATT&CK's chain logic (map

-ping TTP sequences) provides a strong foundation for correlations like the -se.

**Feedback Loops:**

Every false positive alert is also data. Capture analyst feedback automatical -ly and feed it back into rule tuning. Over time, your SIEM "learns" what no -rmal looks like for your environment.

## Step 3: Tune Your EDR/XDR for Precision

We all know Endpoint detection platforms generate massive telemetry. If it is left untuned, they can overwhelm a SOC faster than any SIEM rule. So, f- ollow the measures given below to tune your EDR/XDR.

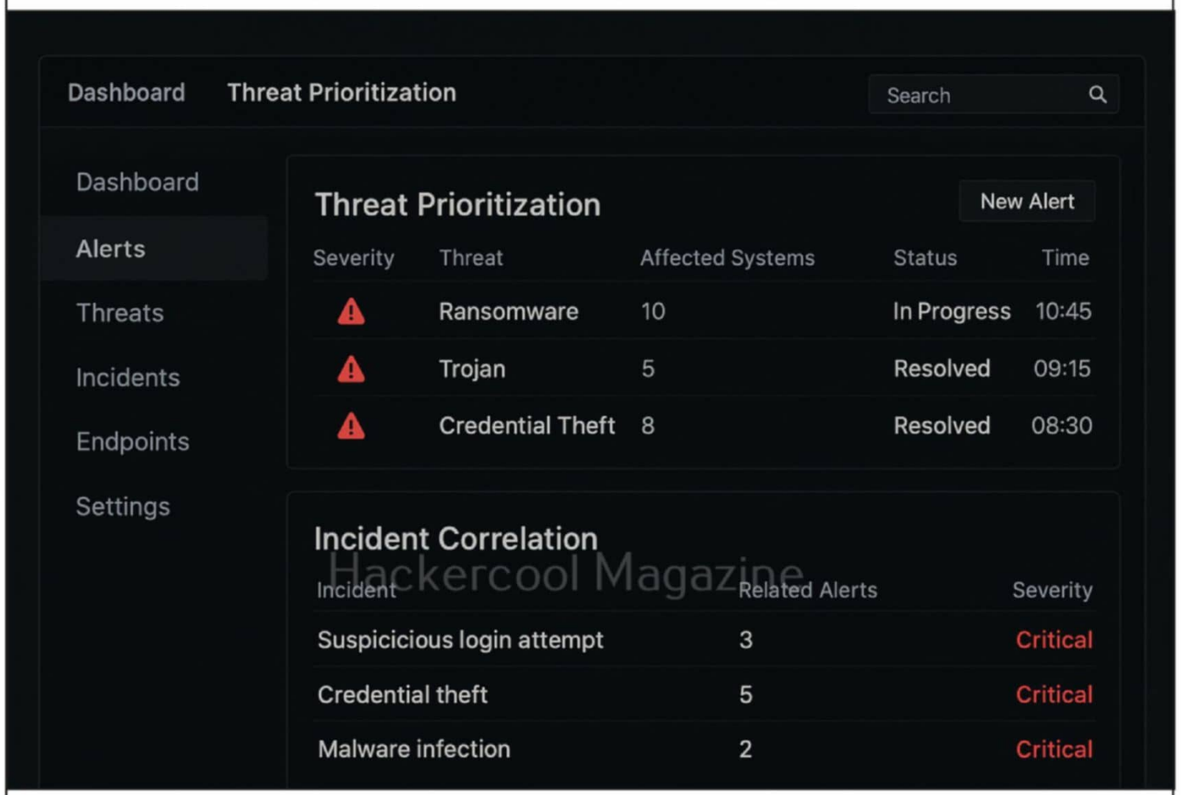


**Focus on Behavioral Detections:**

Static indicators (hashes, filenames) age quickly. So, focusing on behaviora- l detections like suspicious parent-child processes, credential theft patterns, r- egistry persistence etc make it harder for attackers to evade.



**Suppress by Context, Not by Silence:**

Avoid the nuclear option of disabling rules entirely. Instead, use suppressio -n logic:

- Ignore known management tools executing legitimate commands.
- Apply allow-lists for trusted automation.
- Suppress repetitive events that trigger within the same timeframe.

**Use Telemetry Strategically:**

Don't collect everything. Some EDR platforms let you choose what telemet -ry to store versus stream on demand. Keep what's useful for detection and hunting, not every API call.

**Integrate Threat Intel Smartly:**

Match EDR findings against curated threat feeds, but filter for relevance. A -n IOC from a region or sector irrelevant to your organization only adds noi- se.

## Step 4: Automate Intelligently

Automation should not replace analysts totally, it should amplify them. Here are some measures you can take for intelligent automation.

**SOAR (Security Orchestration, Automation, and Response):**

Use automation platforms to handle repetitive triage.

- Automatically enrich IPs with threat intel.
- Quarantine hosts matching high-confidence detections.
- Close duplicate alerts automatically.

The goal isn't to replace human judgment but to preserve it for where it's needed most.

**Event Deduplication and Grouping:**

Combine similar alerts from different systems into a single incident view. This not only reduces alert count but also provides analysts with full context at once.

**Playbook Maturity:**

Start with simple automation (enrichment, tagging). Expand to partial containment and eventually full response for verified patterns like known malware or phishing.

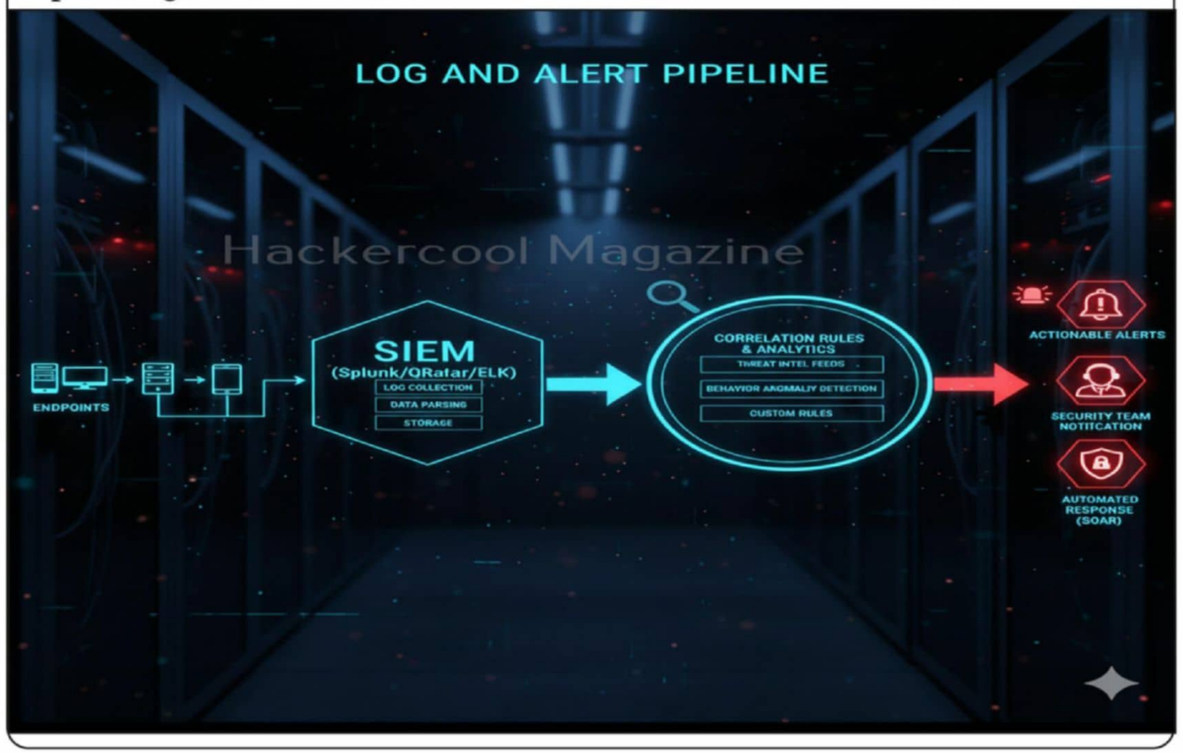




## Step 5: Metrics are important too

Your work is not finished just after implementing all the steps mentioned above. You need to check if our tuning is working as expected. For this observation is needed. Without metrics, tuning becomes guesswork. Track the impact of our noise reduction efforts.

Some key metrics to monitor are,

• True Positive Rate (TPR): This tells us the percentage of alerts that lead to confirmed findings.
• Mean Time to Triage (MTTT): This informs us how long analysts spend v-alidating alerts.
• Alert Volume per Analyst: Trend indicator of load balance.
• False Positive Ratio: This is a core metric for tuning progress.
• Analyst Sentiment Surveys: This should be conducted to check on burno-ut and trust in detections.

If we want to reduce alert fatigue, we should aim for continuous improvem-ent: small, data-driven adjustments that compound over time.

## The Human Element: Designing for the Analyst

Always remember that technology alone can't fix alert fatigue. People proc-ess information differently and cognitive load also matters. So,

**Prioritize UX in the SOC**

If your analysts need five tabs and three consoles to investigate one alert, y-ou've already lost time. Streamline workflows: unify dashboards, automate enrichment and standardize incident templates.

**Empower Analysts to Tune**

The people who triage alerts know what's noise and what's music (I mean i-mportant). Give them authority to propose tuning changes and recognize th-eir contributions. Empowered analysts should feel ownership, not exhaustio-n.

**Build Psychological Resilience**

Encourage micro-breaks, rotations and collaborative reviews in your work-force. The healthiest SOCs should run like sports teams: intense focus, struct-ured practice and real recovery time.

> "Nothing was as satisfying ... than to say that we have run ourselves through the real paces of a real-world ransomware simulation ... mapped as closely as possible to the TTPs of the real adversary."
> –Trevin Edgeworth, Red Team Practice Director at Bishop Fox



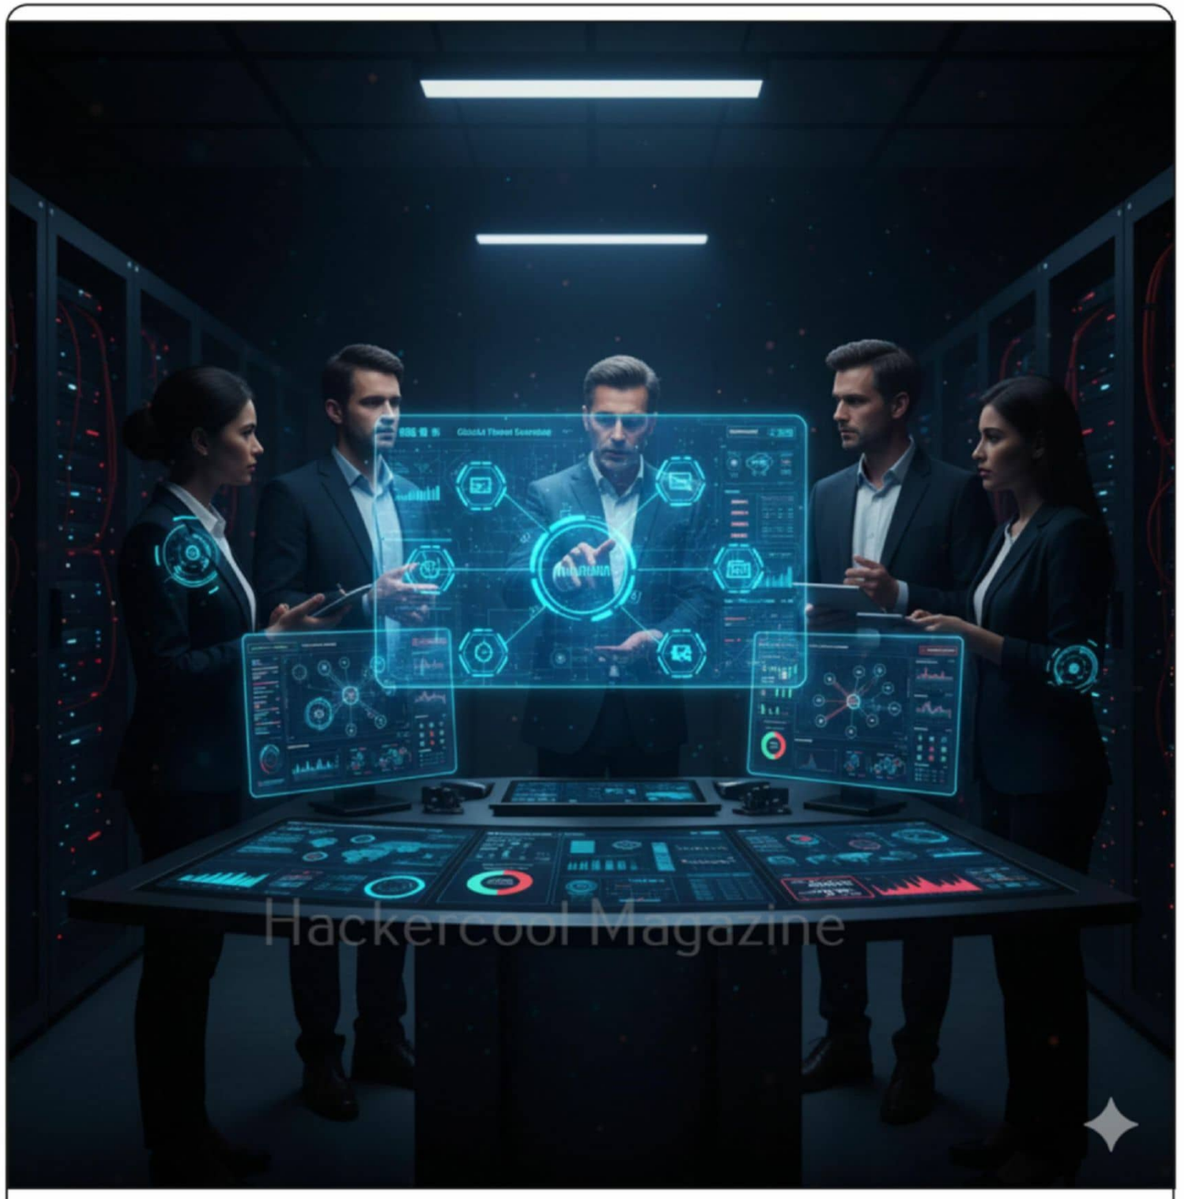


## A Real-World Case Study

At an European organization named LOOKRECKAH, the SOC faced 18,000 daily SIEM alerts. Analysts were regularly missing real threats while drowning in false positives. A Blue Team was made and tasked to reduce th-is alert fatigue. That Team worked for some months and reduced the alert noise by over 80%. How did they do this?

1. They disabled or tuned 40% of default rules based on six months of alert data.
2. Applied MITRE ATT&CK mapping to build 15 correlation chains.
3. Integrated threat intel to auto-enrich and deduplicate incidents.
4. Implemented analyst feedback tracking directly in the SIEM.

Within eight weeks, alert volume dropped by 80% – yet true positive detection increased. Analysts reported lower stress and faster triage. The key here wasn't buying new tools – it was understanding and trusting the ones they already had.

## Continuous Tuning: A Living System

For all its benefits noise reduction is not a single bullet fired once. It's a continuous process. New threats appear. Infrastructure evolves. Detections age. Treat your SIEM and EDR stack like any other piece of software: version control, test cases and iterative improvement. Build a detection backlog. Conduct quarterly tuning sprints. Archive what's obsolete.

Always validate with purple teaming – running simulated attacks to ensure that your tuned rules still catch real adversary behavior.

## Conclusion

In real-world, Blue teams don't win by hearing and responding to every alarm; they win by hearing the right ones and responding to them. A tuned SIEM and EDR ecosystem is the difference between chaos and control – between analysts who chase ghosts and analysts who catch intruders.

The art of noise reduction lies in balance: enough sensitivity to detect true threats, enough precision to ignore the rest.

**Fig: A Blueprint for a Quieter, Smarter SOC**

| Phase | Key Action | Expected Outcome |
|---|---|---|
| Data Audit | Remove redundant or low-value logs | Reduced ingestion cost and noise |
| Rule Tuning | Adjust thresholds, correlation, and enrichment | Higher fidelity alerts |
| Automation | Use SOAR for enrichment and deduplication | Lower manual triage load |
| Feedback Loop | Capture analyst input | Continuous improvement |
| Validation | Red team simulation and metrics review | Confidence in detection coverage |

Hackercool Magazine

"We believe the focus on Samsung Galaxy devices stems from the attackers exploiting a Samsung-specific image-processing zero-day ... we think we're only seeing part of the activity ... this campaign delivering Landfall appears to be part of a broader DNG exploitation wave ... It's also possible ... other mobile vendors were targeted using undiscovered vulnerabilities to deliver the same or similar implants."

- Itay Cohen – Senior Principal Researcher, Unit 42 on Landfall spyware campaign

A critical vulnerability has been detected in Adobe Commerce and Magento Open-Source platform that if exploited can allow hackers to take complete takeover of customer accounts.

Magento®

Hackercool Magazine

## SESSIONREAPER: Session takeover vulnerability

### About the vulnerability

The vulnerability popularly being called as SessionReaper (as it allows hackers takeover entire sessions of customers), carries a CVSS Score of 9.1 out of 10. Tracked as CVE-2025-54236, it is due to an improper input validation.

The vulnerable products are Adobe Commerce (all deployment methods): versions <=2.4.9-alpha2, <=2.4.8-p2, <=2.4.7-p7, <=2.4.6-p12, <=2.4.5-p14 and <=2.4.4-p15. Adobe Commerce B2B versions: <=1.5.3-alpha 2, <=1.5.2-p2, <=1.4.2-p7, <=1.3.4-p14, <=1.3.3-p15. Magento open-source ver -sions: <=2.4.9-alpha2, <=2.4.8-p2, <=2.4.7-p7, <= 2.4.6-p12 and <=2.4.5-p14.

"Magento isn't insecure by design — it's insecure by neglect. The moment you stop patching, attackers start winning."

## How hackers can exploit this vulnerability?

How else do you think hackers exploit this vulnerability? They first scan fo-r Magento stores on internet. There may be over 1,30,000 Magento stores a-round the world but the exact number of vulnerable stores are not known. Once they find the stores running vulnerable versions of Magento, an attack-er crafts input to exploit the vulnerability through its publicly accessible RES-T/web API. Exploitation doesn't require any authentication. If the target Ma-gento store is using file-based session storage, attackers can achieve remote code execution.

## Impact

An attacker can completely take over customer accounts by exploiting this vulnerability. No doubt this vulnerability is being called Session Reaper. Ma-gento instances that do not use file-based session storage (such as Redis bac-



ked) may also be vulnerable.

There have been reports of various incidents where attackers have uploade-d PHP backdoors after exploiting this vulnerability via /customer/address-file/upload as a fake session. In most of the attacks, payloads being used are PHP web shells and phpinfo probes.

It is estimated that by October 26 2025, attackers already targeted over 49% of all Magento stores and 16-18% of stores now have a backdoor installed. M-ost of the attacks are originating from the following IP addresses.

34.227.25[.]4
44.212.43[.]34
54.205.171[.]35
155.117.84[.]134
159.89.12[.]166

## How to manage this vulnerability?

Adobe has released a hotfix to patch this vulnerability during the early day-s of its detection. That was way back in the first week of September 2025. B-ut patching was too slow.

Initially there was no real-world exploitation of this vulnerability. Exploitati-on activity has only picked up recently. At the time of writing this article, ov-er 62% of Magento stores still remain vulnerable to this flaw. This means ov-er 3 in 5 stores ate still vulnerable.

Adobe has advised users to apply the hotfix as soon as possible. Adobe als-o said it released web application firewall (WAF) rules to keep environment-s safe against exploitation attempts.

"It was a precision attack, not a mass campaign ... That strongly suggests espionage motives rather than financial gain."

- Itay Cohen — Senior Principal Researcher, Unit 42 on Landfall spyware campaign

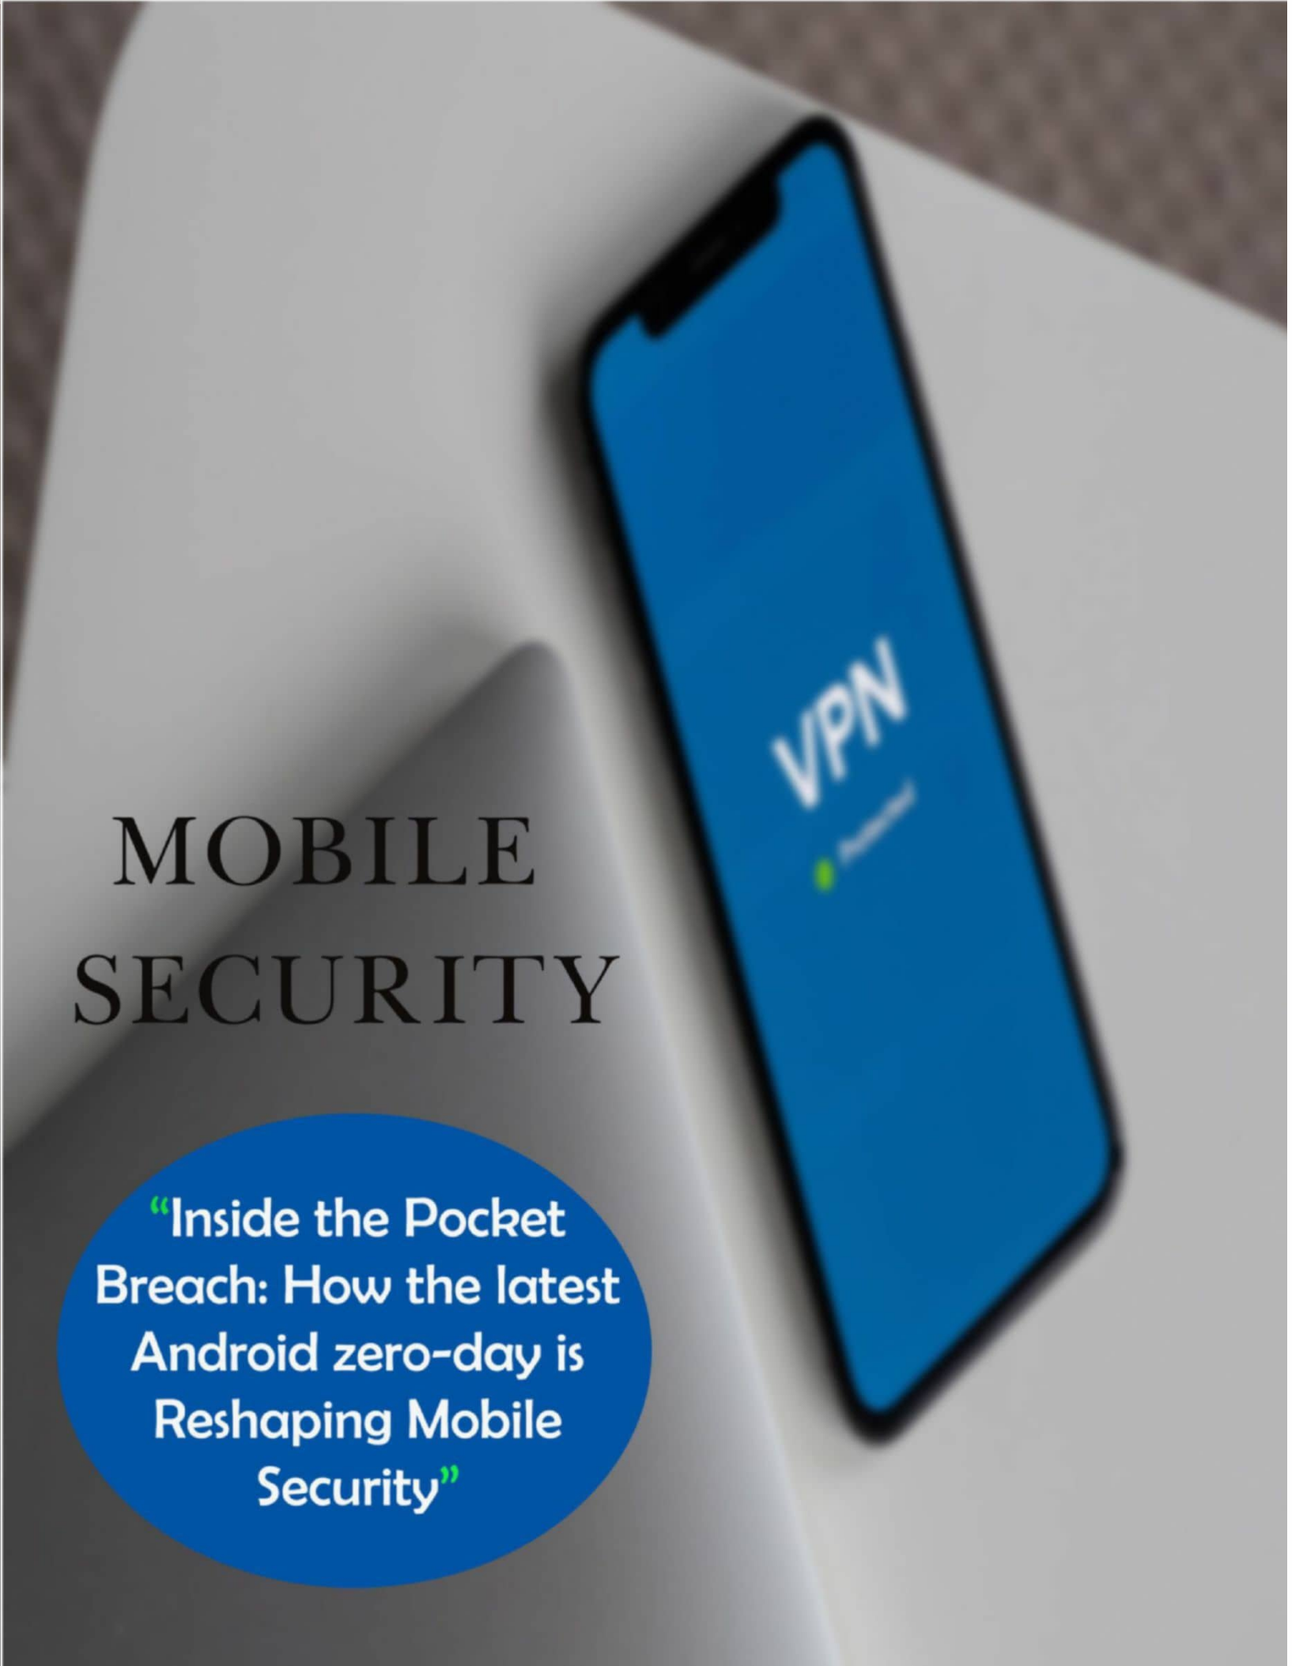
MOBILE
SECURITY
"Inside the Pocket Breach: How the latest Android zero-day is Reshaping Mobile Security"
VPN

Hackercool Magazine

## The Zero-Click Wake-Up Call

In November 2025, the cybersecurity community was jolted by the disclosure of CVE-2025-21042, a zero-day vulnerability in libimagecodec.quram.so, Samsung's proprietary image-processing library. This vulnerability, being exploited in real-world attacks by the LANDFALL spyware campaign, allowed remote code execution with zero user interaction, triggered simply when a malicious image is received and automatically previewed by any messaging app such as WhatsApp.

The implications of this vulnerability are profound: attackers could seize co -ntrol of a device, exfiltrate data and spy on communications without the use -r tapping a single link. For the first time since Pegasus and Hermit, a new generation of mobile sp- yware had entered the battlefield.

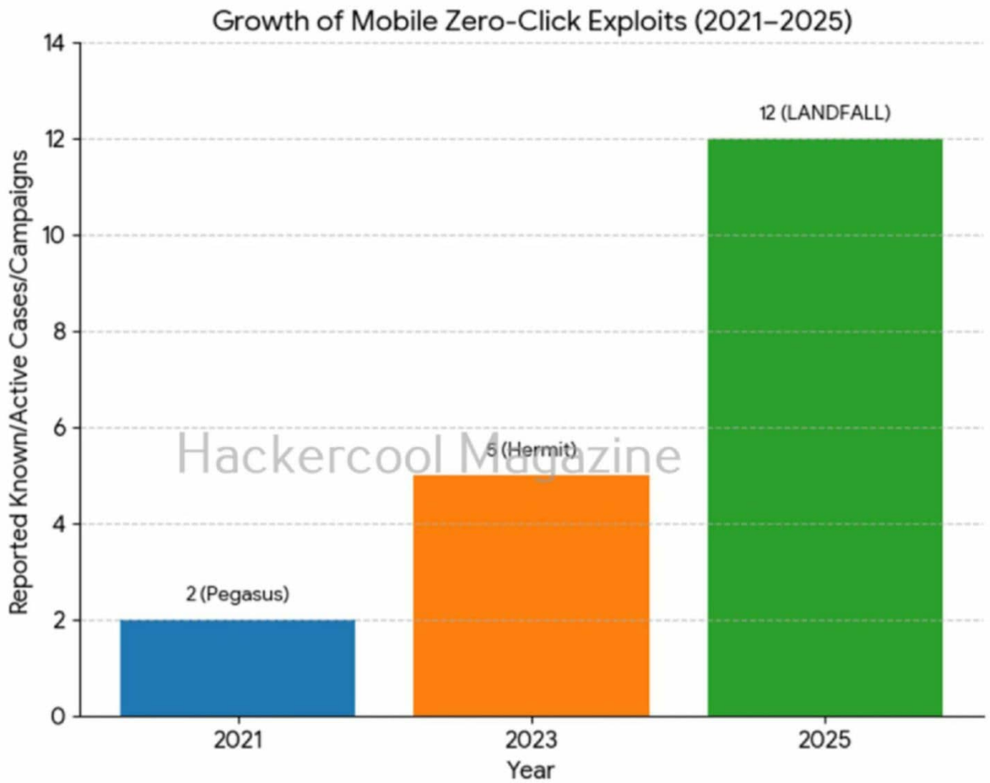




## How the Attack Works

**From Image to Intrusion**

The exploit targets a memory corruption flaw in Samsung's image codec. This codec is responsible for handling DNG image files. DNG stands for Di -gital Negative, which is a raw image file format developed by Adobe that is widely compatible and designed for long-term archiving. It stores a larger a- mount of image data compared to formats like JPEG and is used by professi -onal photographers to maintain image quality and flexibility during editing. Attackers embedded malicious payloads inside carefully crafted images and distributed them via messaging apps. When the recipient's phone just gener- ated a preview thumbnail, the payload executed silently.

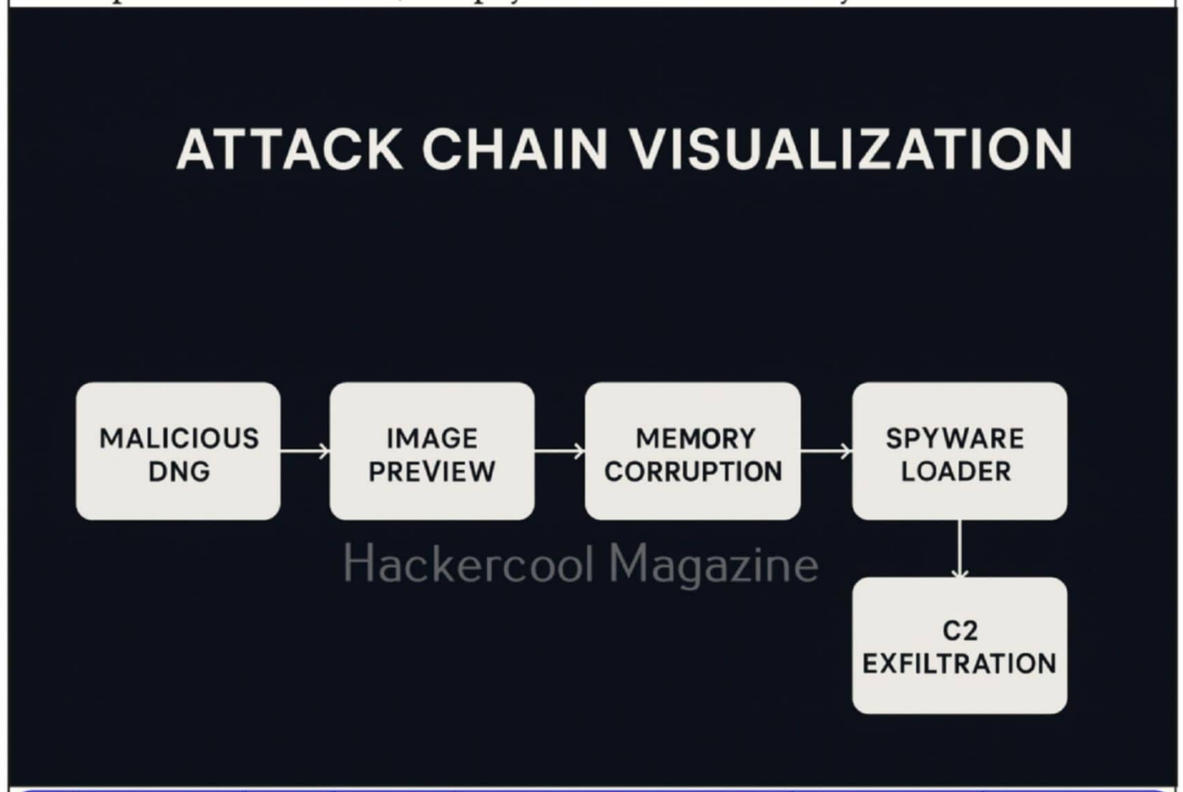


"Logs are the exhaust of your environment; detection is the engine that turns it into intelligence."

## The Attack Chain

**1.Delivery** – A malicious DNG file is sent over WhatsApp or Telegram to the target user.

**2.Trigger** – Image previewing process of Samsung parses the file, causing memory corruption.

**3.Execution** – The attacker's payload executes in the context of the imag-e parser.

**4.Persistence** – Spyware modules get installed (b.so and l.so) and modify SELinux rules. Security-Enhanced Linux (SELinux) is a Linux kernel securit-y module that provides a mechanism for supporting access control security policies, including mandatory access controls (MAC).

**5.Command & Control** – What else, the device connects to remote serve-rs for data exfiltration.

Researchers at Palo Alto Networks' Unit 42 observed that LANDFALL use-d HTTPS for making encrypted communications, modular payload loading and stealth techniques to disable certain logging functions on the device. Samsung phones affected by this vulnerability include Galaxy S22, S23, S24, Z Fold 4, and Z Flip 4, devices frequently used in corporate BYOD environ-ments.

## The Scope of the Campaign

**A Regional Focus — For Now**

Telemetry suggests that early infection activity centered on targets only in t-he Middle East, including Turkey, Morocco, Iraq and Iran. But this may be just a rehearsal. Analysts believe the operation's primary motive was espiona-ge, not cybercrime.

What makes this zero-day especially dangerous is its vendor specificity. Att-ackers rather than exploiting Android's base OS (which we usually see), tar-geted a specific library belonging only to Samsung, highlighting how OEM



customization creates new security threats across the Android ecosystem.-

## Why Mobile Is the New Perimeter?

The LANDFALL campaign reinforces a critical reality: mobile devices are now prime targets for advanced threat actors. Phones are no longer peripher-al tools. They are authentication tokens, communication hubs and data vaul-ts. Compromise a smartphone and you compromise the user's digital identit-y.

The fragmented nature of Android's update cycle only worsens this risk. While Samsung issued a patch in April 2025, thousands of users and entire enterprise fleets remained unpatched months later.

## Defense in Depth for Mobile Devices

**For Enterprises:**

**Patch Enforcement:** Verify that all corporate-issued Galaxy devices have the April 2025 patch or newer.

**Mobile Threat Defense (MTD):** Tools such as Zimperium, Lookout or Microsoft Defender for Endpoint (Mobile) can detect SELinux tampering a-nd rogue outbound traffic.

**Policy Hardening:** Disable automatic image previews in messaging apps where feasible.

**SOC Integration:** Ingest mobile telemetry into SIEM or EDR pipelines for early anomaly detection.

**BYOD Security Controls:** Set minimum patch levels before allowing devi-ce access to corporate resources.

**For Individual Users:**

- Apply Samsung's latest firmware update immediately.
- Avoid opening unsolicited images, especially DNG or RAW formats.
- Keep auto-download and preview options turned off in messaging apps.
- Watch for symptoms such as overheating, camera activation or fast battery drain.

## The Road Ahead

Looking forward, experts expect attackers to increasingly exploit multimed-ia parsing libraries, which handle images, videos and audio files — compon-ents deeply integrated into every smartphone.

As 5G, IoT and wearable tech expand, the attack surface will grow in both scale and complexity.

"Every lock can be picked, given the right key — or the right mind."

— Anonymous



Fig: Emerging Mobile Attack Surfaces — 2026+

| Surface | Risk Type | Example |
|---|---|---|
| AR Glasses | Sensor hijack | Camera feed theft |
| Wearables | Data sync abuse | BLE exfiltration |
| App SDKs | Supply-chain exploit | Malicious SDK updates |

Hackercool Magazine

**Predicted Trends for 2026:**

- **Cross-App Exploits:** Shared media components targeted across multiple messaging platforms.
- **Firmware-Level Attacks:** Exploits chaining chipset and vendor code for persistence.
- **AI-Driven Defenses:** MTD systems using behavioral baselines to detect anomalies.

For defenders, the takeaway is clear — mobile threat visibility must become a core SOC function. Traditional endpoint protection alone is no longer eno-ugh.

## Conclusion

**The Pocket Is Now the Perimeter**

CVE-2025-21042 and the LANDFALL campaign signal a decisive shift in a-ttacker's strategy. The days of "phishing for clicks" are giving way to zero-click infiltration, exploiting features users never interact with.

Enterprises must treat smartphones as first-class endpoints, integrate mobile telemetry into threat detection, and enforce rapid patch compliance.

The true hacker is not a criminal; he is a rebel with a cause — to understand, to question, and to improve."

— Eric S. Raymond

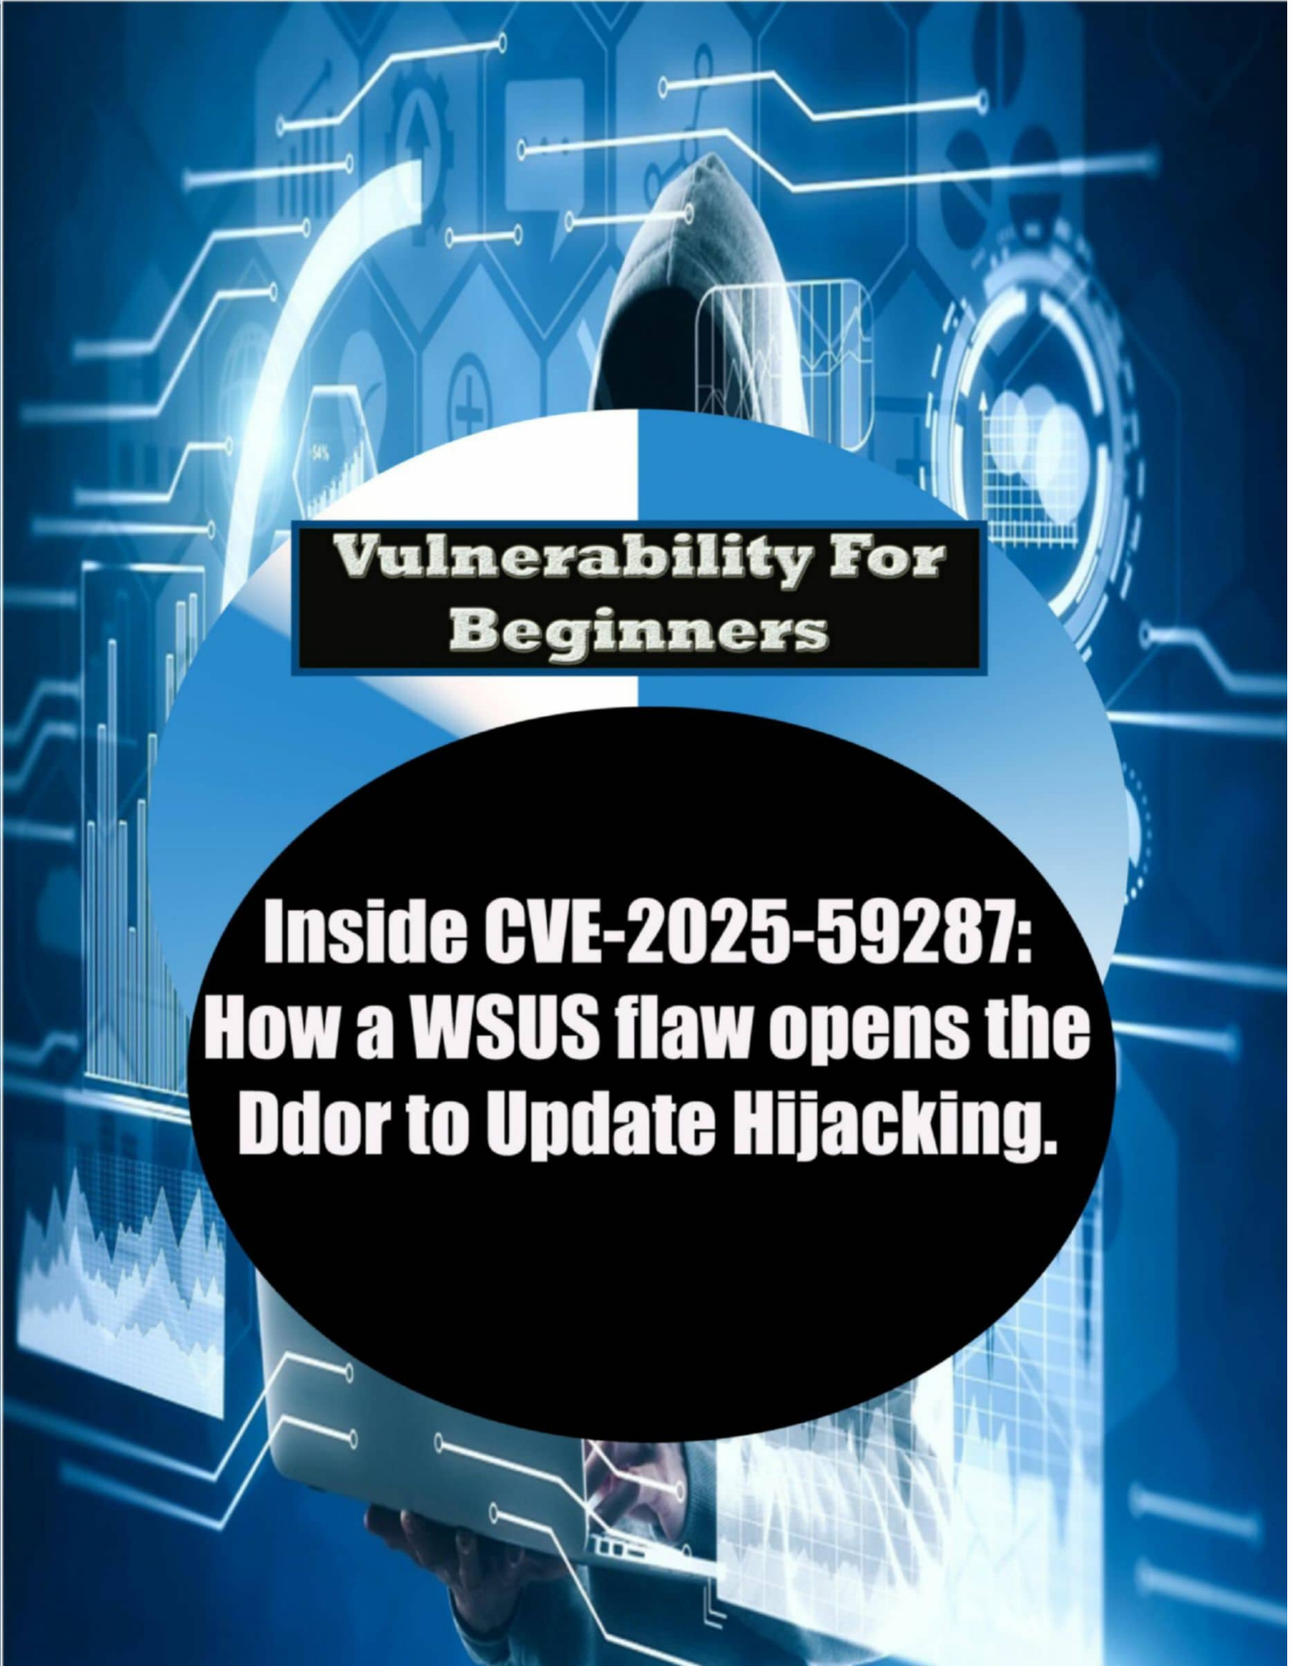


A critica1 vulnerability has been reported in Microsoft Windows Server Updates (WSUS) that allows unauthenticated attackers to execute code remotely on Windows Servers.

**Remote Code Execution vulnerability in Windows Server Update Service (WSUS)**

## About the vulnerability

Tracked as CVE-2025-59287, the vulnerability has a CVSS 3.1 score of 9.8. The vulnerability is due to unsafe deserialization of untrusted data in WSUS's handling of Authorization cookie. Windows Server Update Service (WSUS) is a role in Windows Server that is used by administrators to deploy Microsoft updates across the network.

The vulnerability affects all supported Windows Server versions from 2012 to 2025 where GetCookie's end point processes encrypted Authorization Co -okie () objects without validation.

The deserialization vulnerability is present exactly in EncryptionHelper's DecryptData() method. This method encrypts incoming cookie data with AES-AES-128-CBC cipher and decrypts that data before passing it to NET's Bin-ary Formatter for deserialization.

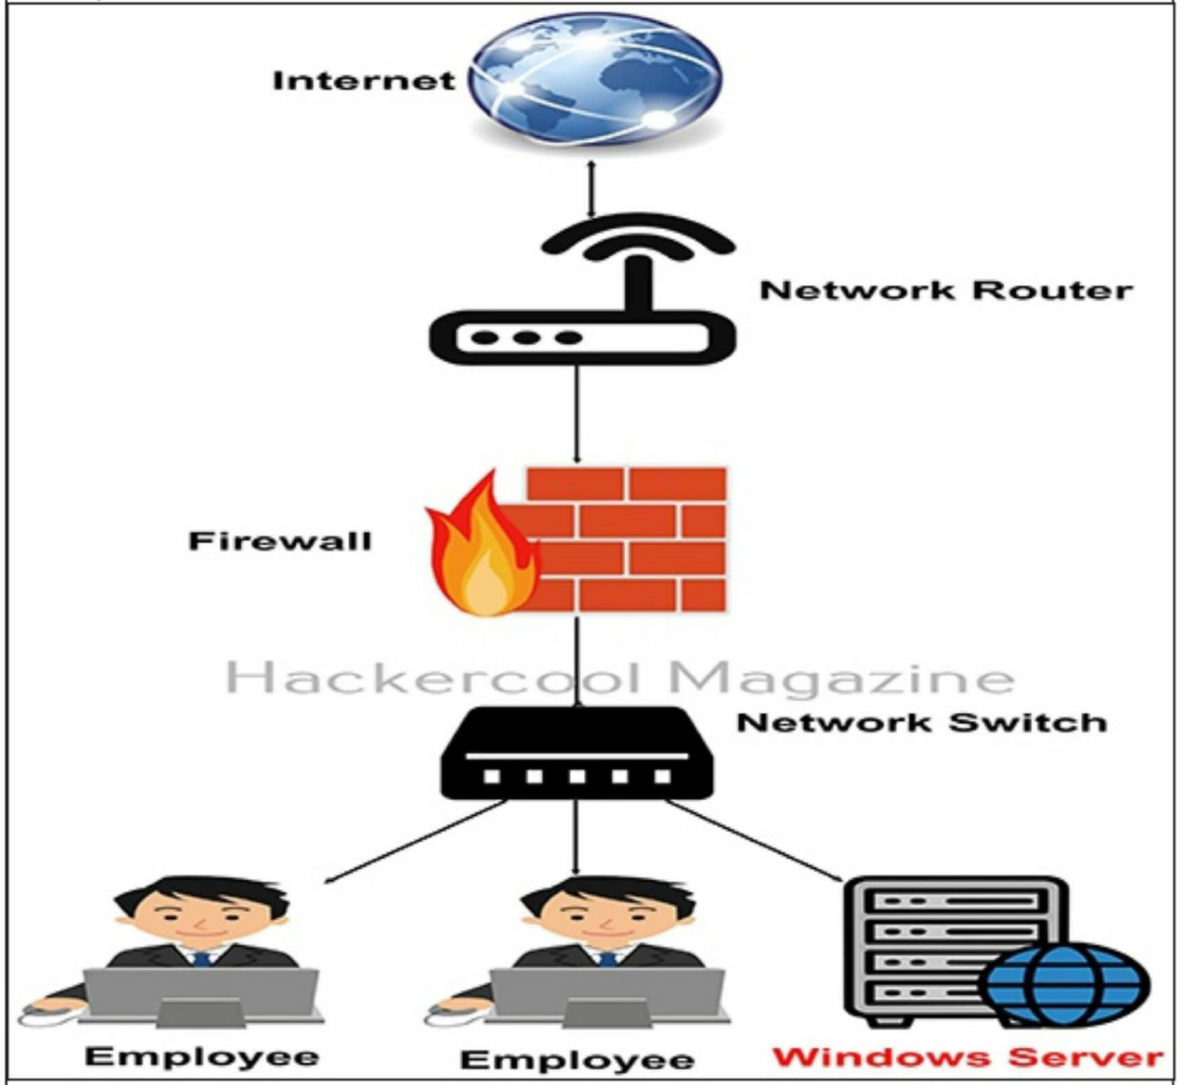


Hackercool Magazine

The problem is this legacy serialization doesn't have any type restrictions and this allows hackers to craft malicious payloads.

## How hackers exploit this vulnerability?

All hackers need to do to exploit this vulnerability is to create SOAP envel-



ope containing a tampered Authorization Cookie and encrypted payload d-ata and send it as an unauthenticated HTTP POST request to the WSUS cli-ent Web Service endpoint running on default port 8530.

The server decrypts the cookie and deserializes the decrypted data via Bin-ary Formatter. The payload gets executed with SYSTEM privileges thus allo-wing RCE.

Hackercool Magazine

Fig: Exploitation scenario of CVE-2025-59287 vulnerability in Windows Server

## Impact

Microsoft classified this vulnerability as "Exploitation more likely" as this v-ulnerability has wormable potential. Exploiting this vulnerability doesn't nee-d any user interaction and this can be dangerous for any exposed WSUS in

-stances. After exploitation, hackers can deploy ransomware, malware or ca-
n compromise the entire network. This can eventually lead to data breaches.

A Proof of Concept (PoC) has been released for this vulnerability and that makes it even more dangerous. Dutch cybersecurity firm Eye Security repor -ted instances of increased activity of scanning of port 8530 and exploitation attempts. WSUS servers are not usually exposed to internet by default but E-ye security found around 2500 instances worldwide.

Even cybersecurity firm Huntress also found evidence of exploitation on d-efault ports of WSUS, 8530 (HTTP), 8531(HTTPS). TrendMicro says it dete-cted over 5,00,000 internet facing WSUS servers.

## How to manage this vulnerability?

Microsoft fixed this vulnerability in October 2025 Patch Tuesday along wit -h 172 other vulnerabilities. Users should update their systems as soon as pos -sible.

These are the updates they should make sure are installed.
Windows Server 2025 (KB5070881),
Windows Server, version 23H2 (KB5070879)
Windows Server 2022 (KB5070884)
Windows Server 2019 (KB5070883)
Windows Server 2016 (KB5070882)
Windows Server 2012 R2 (KB5070886)
Windows Server 2012 (KB5070887)

Once you apply updates, reboot the server for patches to take effects.

Enterprises should also check and disable WSUS service if they are not usi-ng it. If you are using it, isolate the service, Disable WSUS server if not bein-g used, Block inbound traffic at the firewalls to ports 8530 and 853 and chec -k for anomalous SOAP traffic.

As already mentioned, WSUS is a role in Windows Servers that acts as an update service for other WSUS servers in the network. It is not enabled by d -efault.

Uri Gal
Professor in Business Information Systems, University of Sydney

Last week, OpenAI unveiled ChatGPT Atlas, a web browser that promises to revolutionise how we interact with the internet. The company's CEO, Sam Altman, described it as a "once-a-decade opportunity" to rethink how we browse the web.

The promise is compelling: imagine an artificial intelligence (AI) assistant that follows you across every website, remembers your preferences, summarises articles, and handles tedious tasks such as booking flights or ordering groceries on your behalf.

But beneath the glossy marketing lies a more troubling reality. Atlas is designed to be "agentic", able to autonomously navigate websites and take actions in your logged-in accounts. This introduces security and privacy vulnerabilities that most users are unprepared to manage.

While OpenAI touts innovation, it's quietly shifting the burden of safety onto unsuspecting consumers who are being asked to trust an AI with their most sensitive digital decisions.

## What makes agent mode different

At the heart of Atlas's appeal is "agent mode".

Unlike traditional web browsers where you manually navigate the internet, agent mode allows ChatGPT to operate your browser semi-autonomously. For example, when prompted to "find a cocktail bar near you and book a table", it will search, evaluate options, and attempt to make a reservation.

*"While OpenAI touts innovation, it's quietly shifting the burden of safety onto unsuspecting consumers who are being asked to trust an AI with their most sensitive digital decisions."*

The technology works by giving ChatGPT access to your browsing context. It can see every open tab, interact with forms, click buttons and navigate between pages just as you would.

Combined with Atlas's "browser memories" feature, which logs websites

**(Cont'd On Next Page)**



you visit and your activities on them, the AI builds an increasingly detailed understanding of your digital life.

This contextual awareness is what enables agent mode to work. But it's also what makes it dangerously vulnerable.

## A perfect storm of security risks

The risks inherent in this design go beyond conventional browser security concerns.

Consider prompt injection attacks, where malicious websites embed hidden commands that manipulate the AI's behaviour.

Imagine visiting what appears to be a legitimate shopping site. The page, however, contains invisible instructions directing ChatGPT to scrape personal data from all open tabs, such as an active medical portal or a draft email, and then extract the sensitive details without ever needing to access a password.

Similarly, malicious code on one website could potentially influence the AI's behaviour across multiple tabs. For example, a script on a shopping site could trick the AI agent into switching to your open banking tab and submitting a transfer form.

*"Malicious code on one website could potentially influence the AI's behaviour across multiple tabs. For example, a script on a shopping site could trick the AI agent into switching to your open banking tab and submitting a transfer form."*

Atlas's autofill capabilities and form interaction features can become attack vectors. This is especially the case when an AI is making split-second decisions about what information to enter and where to submit it.

The personalisation features compound these risks. Atlas's browser memories create comprehensive profiles of your behaviour: websites you visit, what you search for, what you purchase, and content you read.

While OpenAI promises this data won't train its models by default, Atlas is still storing more highly personal data in one place. This consolidated trove of information represents a honeypot for hackers.

Should OpenAI's business model evolve, it could also become a gold mine for highly targeted advertising.

OpenAI says it has tried to protect users' security and has run thousands

**(Cont'd On Next Page)**

of hours of focused simulated attacks. It also says it has “added safeguards to address new risks that can come from access to logged-in sites and browsing history while taking actions on your behalf”.

However, the company still acknowledges “agents are susceptible to hidden malicious instructions, [which] could lead to stealing data from sites you’re logged into or taking actions you didn’t intend”.

## A downgrade in browser security

*"Atlas is especially vulnerable because it gives human-level control to an intelligence layer that can be manipulated by reading a single malicious line of text on an untrusted site."*

This marks a major escalation in browser security risks.

For example, sandboxing is a security approach designed to keep websites isolated and prevent malicious code from accessing data from other tabs. The modern web depends on this separation.

But in Atlas, the AI agent isn’t malicious code – it’s a trusted user with permission to see and act across all sites. This undermines the core principle of browser isolation.

And while most AI safety concerns have focused on the technology producing inaccurate information, prompt injection is more dangerous. It’s not the AI making a mistake; it’s the AI following a hostile command hidden in the environment.

Atlas is especially vulnerable because it gives human-level control to an intelligence layer that can be manipulated by reading a single malicious line of text on an untrusted site.

## Think twice before using

Before agentic browsing becomes mainstream, we need rigorous third-party security audits from independent researchers who can stress-test Atlas’s defenses against these risks. We need clearer regulatory frameworks that define liability when AI agents make mistakes or get manipulated. And we need OpenAI to prove, not simply promise, that its safeguards can withstand determined attackers.

For people who are considering downloading Atlas, the advice is straightforward: extreme caution.

**(Cont'd On Next Page)**



If you do use Atlas, think twice before you enable agent mode on websites where you handle sensitive information. Treat browser memories as a security liability and disable them unless you have a compelling reason to share your complete browsing history with an AI. Use Atlas’s incognito mode as your default, and remember that every convenience feature is simultaneously a potential vulnerability.

The future of AI-powered browsing may indeed be inevitable, but it should not arrive at the expense of user security. OpenAI’s Atlas asks us to trust that innovation will outpace exploitation. History suggests we shouldn’t be so optimistic.

**This article first appeared in The Conversation**

## DOWNLOADS

Security Onion 2.4.180
**https://docs.securityonion.net/en/2.4/download.html**

Check whether your email is a part of any data breach
**https://haveibeenpwned.com**

**Follow Hackercool Magazine for latest updates**

Buy now

Buy now

Buy now

Buy now

Buy now

Buy now

Buy now

Buy now

2165451